A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY

Military Operations Research

Volume 3 Number 2

1997

A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY
A JOURNAL OF THE MILITARY
OPERATIONS RESEARCH SOCIETY

MORS
MILITARY OPERATIONS RESEARCH SOCIETY

Military Operations Research

A publication of the Military Operations Research Society

The Military Operations Research Society is a professional society incorporated under the laws of Virginia. The Society conducts a classified symposium and several other meetings annually. It publishes proceedings, monographs, a quarterly bulletin, *PHALANX*, and a quarterly journal, *Military Operations Research*, for professional exchange and peer criticism among students, theoreticians, practitioners and users of military operations research. The Society does not make or advocate official policy nor does it attempt to influence the formation of policy. Matters discussed or statements made in the course of MORS symposia or printed in its publications represent the opinions of the authors and not the Society.

Military Operations Research, the journal of the Military Operations Research Society (ISSN 0275-5823) is published quarterly by the Military Operations Research Society, 101 South Whiting Street, Suite 202, Alexandria, VA 22304-3418. The domestic subscription price is $40 for one year and $75 for two years; international rates are $80 for one year and $150 for two years. Periodicals Postage Paid at Alexandria, VA. and additional mailing offices.

POSTMASTER: Send address changes to *Military Operations Research*, the journal of the Military Operations Research Society, 101 South Whiting Street, Suite 202, Alexandria, VA 22304.

Table of Contents

Volume 3, Number 2

1997

CALL FOR PAPERS

SPECIAL ISSUE OF *MILITARY OPERATIONS RESEARCH* ON

WARFARE ANALYSIS AND THE NEW SCIENCES

In September 1997 MORS held a Mini-symposium/Workshop on Warfare Analysis and Complexity. The Mini-symposium drew 165 participants who were enthusiastic about the use of New Sciences in military operations research. The recent article in *PHALANX* summarized much of the discussion regarding the applicability of the New Sciences to problems in military operations research. We want to expand the discussion and consider technical applications of the New Sciences to specific problems in operations analysis. The special issue of *Military Operations Research* will be focused on this topic. The purpose of this special issue is to assess, clarify, and discuss the effectiveness of novel New Sciences techniques and to propose improvements to military operations analysis.

The New Sciences are meant to include chaos and complexity theory, edge of chaos, complex adaptive systems, and more generally, any techniques for rapidly scanning the solution space of a computationally complex problem in order to determine boundaries of behavior and regions where rapid changes occur. All papers will be unclassified. In accordance with *MOR* editorial policy, we require certification from a senior decision-maker when applicable.

Interested authors should submit ABSTRACTS by **31 May 1998** and PAPERS by **30 September 1998**. The papers should be submitted in accordance with our current editorial policy. All papers will be refereed.

Please contact me if you are interested in authoring a paper for this special issue.

DR. JULIAN PALMORE
Guest Editor, *Military Operations Research*
University of Illinois at Urbana-Champaign
1409 W. Green Street
Urbana IL 61801-2975
Phone: 217-333-0407 (with voice mail)
Email: palmore@uiuc.edu

AN INTEGER SOLUTION HEURISTIC FOR THE ARSENAL EXCHANGE MODEL (AEM)

by Daniel J. Green, James T. Moore, and John J. Borsi

Air Force Studies and Analyses Agency (AFSAA) was using the Arsenal Exchange Model (AEM) to allocate the weapons of various strategic force structures to targets. AEM often allocates a fractional number of weapons to a class of targets and then truncates this number to get an integer allocation of weapons. AFSAA wanted a better approach to obtain integer solutions. In this research, Dan Green, Jim Moore, and John Borsi developed a heuristic which finds better integer solutions in less than five seconds. Their heuristic has been incorporated in AEM.

ON THE PRINCIPLE OF "FALSE RANGING" IN WEAPONRY

by Israel David

Suppose you are an army commander who is requested to join a multinational force, complying with others' guns or ammunition. Or suppose a serious "bug" has been detected in a crucial weapon fire control system. Or imagine any other unpredicted situation that arises which calls for an unconventional use of weapons. Is there any methodology that may help in suggesting a relatively handy, simple and accurate means for coping with such problems in times of emergency?—The present paper further clarifies, defines and exemplifies the problem. It discusses the apparently old approach to its solution in simple contexts, coined here "false ranging." Moreover, it shows how modern OR techniques may be used to generalize the principle and to enable its application in larger contexts.

RISKNAV™: A DECISION AID FOR PRIORITIZING, DISPLAYING, AND TRACKING PROGRAM RISK

by C. C. Cho, P. R. Garvey, and R. J. Giallombardo

This work presents a family of preference models for prioritizing program risk. These models originate from multiattribute utility theory and rank-order project-defined risk events as a function of multiple criteria. Such criteria includes, but is not limited to, a program's cost, schedule, and technical performance. In addition, the methodology is tuned for quantifying the effects of coupled (dependent) risks. As a decision-aid, these models target where engineering assets are best applied to mitigate potentially crippling areas of risk to a program.

THE PROBABILISTIC MULTIPLE-TRAVELLING-SALESMEN FACILITY-LOCATION PROBLEM: SPACE-FILLING CURVES AND ASYMPTOTIC EUCLIDEAN ANALYSES

by Yupo Chan and David L. Merrill

All of us need to make decision in uncertainty, particularly regarding the demands that are placed upon our day-to-day operations. The operations have to be executed in real time, and long-term plans have to be made. In the defense community, we face such issues as base closures and fleet requirements to sustain a required mission. When locating the base for a fleet of fligh-inspection aircraft, Yupo Chan and David Merrill show how some analysts in the U. S. Air Force went through these decisions. While flight-inspection requirements vary from day to day, the authors offer a robust analysis procedure that can respond to these stochastic demands. Furthermore, the same analysis procedure, when carried out over a sufficient lengthy period of time (such as a year), can suggest basing and fleet size alternatives. Best of all, the procedure can be executed in the field with minimal computational requirements. It offers a fast solution to a highly taxing technical problem in Operations Research, while guaranteeing an error bound to the "quick solution" (a "75 percent solution"). The model reflects a current trend in combinatorial optimization, wherein the best of heuristics with analytical formulation are combined in a single model.

APPLICATION AND EXTENSION OF THE THRUPUT II OPTIMIZATION MODEL FOR AIRLIFT MOBILITY

by Richard E. Rosenthal, Steven F. Baker, Lim Teo Weng, David F. Fuller, David Goggins, Ayhan 0. Toy, Yasin Turker, David Horton, Daniel Briand, and David P. Morton

Recent advances in computing and model formulation have permitted optimi-

Executive Summaries

zation models to influence airlift mobility decision making. This paper describes the use of a large-scale linear program to provide insight to the C-17 Defense Acquisition Board. For this analysis, we used a variety of airlift fleets under a two Major Regional Contingency scenario. We also describe ongoing efforts to improve the realism and tractability of optimization models for airlift anayses, including model reduction and decomposition techniques, as well as incorporating the stochastic nature of aircraft breakdowns.

An Integer Solution Heuristic for the Arsenal Exchange Model (AEM)[‡§]

Daniel J. Green,
James T. Moore,

AFIT/ENS,2950 P Street, Wright-Patterson AFB, OH 45433-7765

John J. Borsi

AMCSAF, 402 Scott Dr, Unit 3L3, Scott AFB, IL 62225-5307

ABSTRACT

The Air Force Studies and Analyses Agency (AFSAA) uses the Arsenal Exchange Model (AEM) to allocate weapons to targets. AFSAA needed an improved method for converting the continuous solution produced by AEM into a feasible integer solution. The current integerization method leaves weapons unused and targets uncovered. In the method developed in our research, the noninteger valued variables in the continuous solution are truncated to produce an integer solution. An integer goal program is implemented to reallocate the weapons and targets made available in the truncation process. The truncated solution is then combined with the results of the integer goal program to produce a feasible integer solution for the original problem. Our method, using the linear programming relaxation of the integer goal program, was implemented in AEM. The implementation was used to solve four test cases. In all four cases, our method produced results that were closer to the continuous solution in terms of damage expectancy, target coverage, and goal performance than the results produced by the current method.

INTRODUCTION

Since the dawn of the nuclear age, the United States' defense community has been analyzing the capabilities of the nuclear arsenals of both the United States and its potential adversaries. One approach to this analysis is to assign the weapons of one country to targets in another country. When a weapon is assigned to a target, this means the specified weapon will be used to attack the designated target. To measure the damage achieved by this assignment, the weapon's probability of arrival at the target and the weapon's probability of damaging the target are used to compute damage expectancy.

The number of weapons can exceed 10,000 and the number of potential targets can also exceed 10,000. Thus, the problem of assigning weapons to targets can be very large with over 100 million variables representing possible weapon to target assignments. The problem is made more challenging by the fact that the number of weapons assigned to a particular target can be greater than one.

The Arsenal Exchange Model (AEM) represents this assignment problem by grouping similar weapons and similar targets. The available weapons are divided into types. Examples of weapon types are B-52 bombs, B-52 cruise missiles, Minuteman III warheads, and Trident D-5 warheads. The targets are grouped by class. Examples of target classes are tank factories, submarine pens, missile silos, and airfields. For each weapon type, the number of available weapons and the weapon's characteristics are modeled. Weapon characteristics include: yield, reliability, and accuracy. For each of the target classes, the number of targets and their characteristics are modeled. Target characteristics include: hardness, radius, and value (Bozovich et al. 1993:149-151).

Using this grouping of weapons and targets, AEM allocates weapons to targets for either side in a scenario (Bozovich et al. 1993:4). The allocation's measure of goodness is damage expectancy (DE). DE for a particular weapon-target combination is the product of a weapon's probability of kill against a target times the target's value. AEM calculates this DE for each possible weapon-target combination and then does an allocation of weapons to targets so as to maximize the sum of damage expectancies. These weapon-target combinations are called strategies, and they are the decision variables in AEM.

An allocation is completely defined by specifying the number of times each strategy is used (Bozovich et al. 1993:154-157). In addition to maximizing damage expectancy, other objectives or requirements for the allocation may be established. Examples of these include a requirement that all targets have a weapon allocated to them, that an average DE of 0.75 be achieved against tank factories, and that each missile silo should be allocated two weapons. These requirements may be added as goals which the allocation should satisfy as closely as possible or as "hard" constraints which the allocation must satisfy. While the hard constraints must be satisfied, even at the expense of reduced DE, the goals do not have to be satisfied. Although AEM tries to satisfy all goals, the resources required to satisfy all goals may not be available. In this situation, AEM gets as close to goal satisfaction as possible. AEM will sacrifice DE to achieve the goals, but when sufficient resources are not available to satisfy all goals, it continues with the solution pro-

APPLICATION AREA:
Weapon Allocation

OR METHODOLOGIES:
Linear Programming, Integer Programming

‡§Approved for Public Release; Distribution Unlimited

cess. Therefore, failure to satisfy a goal affects the goodness of the allocation while failure to satisfy a constraint affects the feasibility of the allocation.

AEM employs a preemptive goal programming approach to solve the problem of allocating weapons to targets when prioritized goals are specified. In AEM, solving the problem of allocating weapons to targets poses two problems. The first problem concerns redundant constraints. AEM adjusts the levels of the goals to avoid the singularity problems redundant constraints cause. The second problem occurs when a goal cannot be satisfied. When this is the case, AEM fixes the goal at the value achieved before proceeding with the inclusion of the next level of prioritized goals. This approach yields a continuous solution of the weapons to target allocation problem.

In the procedure presented in this paper, we develop a superior method for obtaining feasible integer solutions from the continuous solution provided by AEM. Since AEM is also the main solution engine in the Conventional Target Evaluation Methodology (CTEM), our procedure also has potential to improve the performance of this model. CTEM (Gallagher and Kelly 1991) uses AEM to allocate conventional and/or nuclear weapons to a conventional target base. Since an integer solution to this problem aids planning and evaluation, our procedure has the potential to improve CTEM's usefulness and performance.

Because an integer number of weapons must be allocated to an integer number of targets, the problem AEM seeks to solve is an integer programming (IP) problem. However, the problem is so large, often involving more than a million general integer variables, that the time required to find an optimal integer solution is prohibitive in the setting in which AEM is used. Fortunately, defense analysts do not require the optimal integer solution; rather, they require a good feasible integer solution quickly. In order to provide a feasible integer solution in a reasonable amount of time, AEM relaxes the integer requirements and solves the resulting linear programming (LP) problem. AEM then takes the noninteger variables in the continuous LP solution and truncates them. Since this truncation process leaves weapons unused and targets uncovered, it seems likely that the method could be improved (Cotsworth 1993). The objective of this research is to develop a method for obtaining integer solutions for the weapon to target allocation problem which is superior to the current truncation approach. In the following sections, we present the mathematical model used by AEM, AEM's application of goal programming, our solution heuristic and its implementation, the results of testing our heuristic, and some conclusions and recommendations.

AEM'S MATHEMATICAL MODELING

AEM solves a mathematical model to find a feasible allocation of weapons to targets which maximizes DE. This model has the following formulation:

Objective Function

$$Maximize \sum_{i=1}^{m} DE_i \, x_i \tag{1}$$

subject to:

$$\sum_{i \in ST_j} x_i \leq T_j \quad j = 1,2,\cdots,n \tag{2}$$

$$\sum_{i \in SW_k} B_{ki} \, x_i \leq W_k \quad k = 1,2,\cdots,p \tag{3}$$

$$x_i \geq 0 \quad i = 1,2,\cdots,m \tag{4}$$

where:

- m = The number of allowable strategies.
- n = The number of target classes.
- p = The number of weapon types.
- B_{ki} = The number of type k weapons expended by one use of strategy i.
- DE_i = The damage expectancy resulting from one use of strategy i.
- W_k = The number of type k weapons.
- SW_k = The set of strategies which use weapon type k.
- T_j = The number of class j targets.
- ST_j = The set of strategies which attack target class j.
- x_i = The number of times strategy i is used in the allocation.

The objective function (1) maximizes the sum of the damage expectancies of the strategies. For each strategy, the damage expectancy is computed. This computation does not de-

pend on or consider the damage achieved by other strategies. Although this approach does not consider the damage done to other targets by a strategy, this situation can be mitigated by the use of designated ground zeros which group targets together for damage calculation purposes based on the selected weapon. Constraints (2) are the target constraints. They ensure that, for each target class, the number of targets attacked by the strategies is no greater than the number of targets. Constraints (3) are the weapon constraints. They ensure that, for each weapon type, the number of weapons used in the strategies is no greater than the number of weapons available. To this basic formulation, AEM adds the constraints and goals designated by the analyst. Goals specified by the analyst are added to the basic formulation using a goal programming approach.

GOAL PROGRAMMING IN AEM

In addition to maximizing DE when allocating weapons to targets, other goals for the allocation may be established. In AEM, these goals are called hedges. There are four types of hedges (Bozovich et al. 1993:138):

1. Requirements pertaining to the total damage achieved on a specified set of target classes by a specified set of weapon types. An example is allocate weapon types 2 and 3 to achieve an average damage expectancy of at least 0.8 against class A targets.
2. Requirements pertaining to the total number of allocatable weapons from a specified set of weapon types allocated against a specified set of target classes. An example is use no more than one-half of the available type 2 weapons against class A targets.
3. Requirements pertaining to the total number of targets hit by a specified set of weapon types. An example is attack at least one-half of the class A targets using type 2 weapons.
4. Requirements that all strategies for the specified weapon/target combinations satisfy a set of criteria. The criteria deal with the amount of damage expectancy achieved by the strategy, the number of weapons involved, and the presence of certain weapon types in the strategy. An example is attack class A targets using strategies with damage expectancies of at least 0.8.

Goals are often similar to constraints in the problem. For example, a goal may be to allocate a weapon to every class A target. This goal may be written as:

$$\sum_{i \in S_A} x_i + d^- = T_A$$

where S_A is the set of strategies that attack target class A, T_A is the number of class A targets, and d^- is the deviation variable which represents the number of class A targets which do not have a weapon allocated to them. The goal is to minimize the deviation variable d^-. The target constraint for class A targets is:

$$\sum_{i \in S_A} x_i + s = T_A$$

where s is the slack variable for this constraint. If the goal is met, then d^- and s will both be zero and AEM will remove the columns associated with these variables from the tableau. Once the variables d^- and s have been removed, the rows associated with the goal and with the constraint may both be written as:

$$\sum_{i \in S_A} x_i = T_A$$

Duplicate or redundant rows in the tableau result in a singular constraint matrix which causes difficulties for the LP solver used in AEM. Currently, the AEM solver uses a combination of generalized upper bounding (GUB), decomposition, Gauss-Jordan elimination, and compact tableau storage. For large, complex AEM applications, precision and solution time problems are noticeable (Gallagher and Kelly 1991:881). The AEM LP solver uses several techniques to alleviate the precision problems. Among these are reinversion and backward pivoting. Although these techniques alleviate precision problems, those caused by redundant constraints are not entirely corrected. The identification and elimination of redundant constraints is a difficult and challenging problem. One approach to this problem would require the pairwise comparison of the goal and hedge constraints with the weapon and target constraints. For large problems, the amount of time required to accomplish all comparisons would be prohibitive.

As shown above, the redundant constraints are usually caused by goals or hedges. These redundant constraints cause singularity problems in the Gauss-Jordan elimination routine. Since singular matrices cannot be inverted,

AEM multiplies the right-hand-side of all goals by 0.9995. This causes the goals and constraints to differ, but does not change the goals enough to cause a significant change in the solution. That is, the goal performance should only change by 0.05 percent (Cotsworth 1993).

AEM incorporates these goals in the LP relaxation using preemptive goal programming where a series of linear programming subproblems are solved. The first subproblem may be formulated as:

Objective Function

$$Maximize \sum_{i=1}^{m} DE_i\, x_i - \sum_{h \in P^1} M(w_h^+\, d_h^+ + w_h^-\, d_h^-) \quad (5)$$

subject to:

$$\sum_{i \in ST_j} x_i \le T_j \quad j = 1,2,\cdots,n \quad (6)$$

$$\sum_{i \in SW_k} B_{ki}\, x_i \le W_k \quad k = 1,2,\cdots,p \quad (7)$$

$$\sum_{i=1}^{m} a_{hi} x_i + d_h^- - d_h^+ = 0.9995 b_h \quad h = 1,\cdots,q \quad (8)$$

$$x_i \ge 0 \quad i = 1,2,\cdots,m \quad (9)$$

$$d_h^-, d_h^+ \ge 0 \quad h = 1,2,\cdots,q \quad (10)$$

where:

- m = The number of allowable strategies.
- n = The number of target classes.
- p = The number of weapon types.
- q = The number of goals for the allocation.
- B_{ki} = The number of type k weapons expended by one use of strategy i.
- DE_i = The damage expectancy resulting from one use of strategy i.
- W_k = The number of type k weapons.
- SW_k = The set of strategies which use weapon type k.
- T_j = The number of class j targets.
- ST_j = The set of strategies which attack target class j.
- a_{hi} = The coefficient associated with strategy i in goal h.
- b_h = The right-hand-side of goal h.
- w_h^+ = 0, if goal h is a ≥goal, and 1 otherwise.
- w_h^- = 0, if goal h is a ≤goal, and 1 otherwise.
- d_h^+ = The positive deviation variable associated with goal h.
- d_h^- = The negative deviation variable associated with goal h.
- P1 = The set of goals assigned to priority one.
- x_i = The number of times strategy i is used in the allocation.
- M = A very large positive integer.

The objective function (5) maximizes the sum of the damage expectancies of the strategies while imposing a penalty for not satisfying goals. Constraints (6) are the target constraints. They ensure that, for each target class, the number of targets attacked by the strategies is no greater than the number of targets. Constraints (7) are the weapon constraints. They ensure that, for each weapon type, the number of weapons used in the strategies is no greater than the number of weapons available. Constraints (8) incorporate the goals into the mathematical model. Note the right-hand-side is multiplied by 0.9995 to prevent the inclusion of a redundant constraint.

To solve this goal program, AEM uses an enhanced preemptive goal programming approach to allocate weapons to targets so as to maximize damage to targets while achieving specified goals (Gallagher and Kelly 1991:880). As a result, a continuous solution is produced by solving a series of linear programming subproblems. Priorities are assigned to the goals and the subproblems are solved to minimize the deviation from the goals in order of priority. The objective of each subproblem is to maximize DE while minimizing the deviation from the goals at the current priority level without increasing the deviations from the higher priority goals. In the optimal tableau for the above subproblem, if a priority one goal has been met, the weighted deviation variables associated with the goal have a value of zero. If a priority one goal cannot be met, a weighted deviation variable associated with the goal will be non-

zero in the optimal tableau. In this case, the right-hand-side of the goal is adjusted so as to make the deviation variable have a value of zero. Once the weighted deviation variables have been set to zero, they are removed from the tableau. This prevents the solutions of subsequent subproblems from deviating from the level of goal achievement reached in the current subproblem. The weighted deviation variables associated with the priority two goals are then added to the objective function and the new subproblem is solved. The process is continued until the deviation variables of all the goals have been added to the objective function (Cotsworth and Garrett 1991:Sec II, 32-33). The flow of the subproblems is shown in Figure 1. For example, assume the highest priority goal requires the average damage expectancy of class A targets to be at least 0.8. If it is possible to achieve this level of damage expectancy,

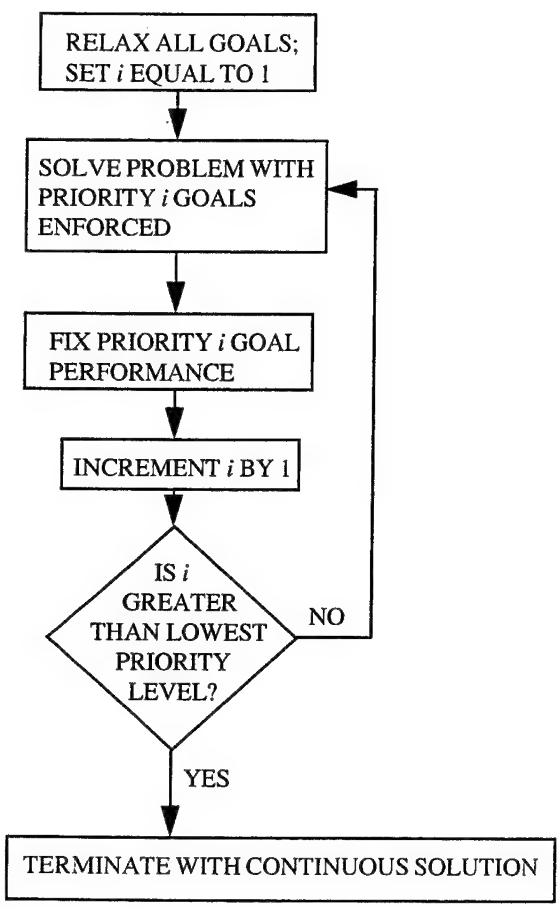

Figure 1. AEM's Preemptive Goal Programming Approach.

then all subsequent subproblems are constrained to also produce an average damage expectancy of at least 0.8. If the greatest possible average damage expectancy against class A targets is only 0.79, then all subsequent subproblems are constrained to produce an average damage expectancy of 0.79 against class A targets.

If the integer restriction on the variables was enforced within the goal programming solution process, significant differences between the integer and continuous solutions could be observed. To see how the solution obtained using LP relaxation would differ from the solution found solving integer programming subproblems, again assume that the highest priority goal is to achieve an average damage expectancy of 0.8 against class A targets. If this goal cannot be met, then the integer solution will have a deviation which is at least as large as the deviation produced by the LP relaxation. For example, the integer solution may achieve an average damage expectancy of 0.78, while the LP relaxation may achieve 0.79. In this situation, subsequent subproblems will be less constrained in the integer case; the integer constrained subproblems require an average damage expectancy of 0.78 while, for the relaxed case, 0.79 is required. Although the integer solution has not done as well as the continuous solution on this goal, the remaining subproblems are less constrained in the integer case, and therefore, the integer solution may perform better than the continuous solution on lower priority goals.

On the other hand, if the goal can be met, then, in the relaxed case, all subsequent subproblems are also required to achieve an average damage expectancy against class A targets of at least 0.8. However, in the integer case, it may not be possible to achieve a damage expectancy of exactly 0.8. It may be that the lowest average damage expectancy which can be achieved by an integer solution, while still achieving at least 0.8, is 0.81. This means that subsequent subproblems are more constrained in the integer case; they must produce an average damage expectancy of at least 0.81 while, in the relaxed case, only 0.8 is required. Although the integer solution may meet this goal at a higher level than the continuous solution, the remaining subproblems are more constrained, and therefore, the integer solution may do worse than the continuous solution on lower priority goals.

Because of the complex balancing which occurs between the goals, it is difficult to predict how the optimal integer solution would compare to the continuous solution. The integer solution may do worse than the continuous solution on some goals and better on others. In addition, because the final subproblem may be more or less constrained, the integer solution may have a damage expectancy which is higher, lower, or exactly the same as the damage expectancy achieved by the continuous solution.

FEASIBLE INTEGER SOLUTIONS

AEM currently produces feasible integer solutions by starting with the solution of the LP relaxation and truncating all of the noninteger valued variables. Since there are an integer number of weapons and targets, many of the variables in the LP solution have integer values. However, when variables which do not have integer values are truncated, the result is unused weapons and uncovered targets. Thus, the truncation can result in an inferior feasible integer solution with reduced damage expectancy.

An examination of the linear programming relaxation of the Generalized Assignment Program (GAP) suggests an approach for achieving better feasible integer solutions for the weapon allocation problem. Benders and van Nunen show that the number of non-unique assignments is less than or equal to the number of machines which are used to capacity. This suggests that the solution of the linear programming relaxation of the GAP (LGAP) may be a good starting point for a heuristic (Benders and van Nunen 1983:48).

One heuristic which builds upon the solution of the LGAP is Trick's LR-Heuristic. In this heuristic, the unsplit jobs are fixed, leaving only the split jobs to be scheduled. Variables associated with machine and job assignments are deleted from the problem if the remaining unused capacity of a machine is not sufficient to perform a job. The resulting LGAP is solved and, if necessary, the procedure is repeated until the integer requirements are satisfied. Trick shows that the procedure will have to be repeated at most m times where m is the number of machines (Trick 1992:140).

Since the problem AEM is designed to solve is different from the GAP, Trick's LR-Heuristic will not be directly applicable to AEM solutions. However, the heuristic developed in this research is similar to Trick's in that it begins with the solution of the LP relaxation, the integer portion of the solution is fixed, and a new problem is formulated to reassign the resources associated with the noninteger portions of the continuous solution.

If, instead of truncating all the noninteger variables, some were rounded up, the damage expectancy and target coverage would be improved, as compared to the current solution method, and the goal performance would still be about the same as in the continuous solution. In fact, since most goals require that at least a certain level of damage expectancy or target coverage be attained, and since rounding some variables up will increase both damage expectancy and target coverage, goal performance should generally be improved compared to the integer solutions obtained using the current method. Also, the objective of the continuous solution produced by AEM is to maximize damage expectancy while satisfying goals which include coverage of important target classes. These observations lead to the use of a similar goal programming approach to ensure coverage of the important target classes while determining which noninteger valued variables in the continuous solution should be rounded up. The heuristic which implements this approach is presented in the next section.

FEASIBLE INTEGER SOLUTION HEURISTIC

Relying on the observations presented in the previous section, an integer goal program may be constructed which maximizes damage expectancy while allocating weapons to every target in those target classes which were completely covered in the continuous solution. Completely covered means the number of strategies allocated to a target class equals the number of targets in that class. This allocation problem is called the Rounding IP and may be formulated as follows:

Objective Function

$$Maximize \sum_{i=1}^{m} DE_i x_i - Md^- \qquad (11)$$

subject to:

$$\sum_{i \in ST_j} x_i \leq T_j \quad j = 1, \cdots, n \tag{12}$$

$$\sum_{i \in SW_k} B_{ki} x_i \leq W_k \quad k = 1, \cdots, p \tag{13}$$

$$\sum_{i \in SC} x_i + d^- = \sum_{j \in TC} T_j \tag{14}$$

$$x_i \in (0,1) \; i = 1, \cdots, m \tag{15}$$

$$d^- \geq 0 \tag{16}$$

where:

- m = The number of strategies with noninteger value in the continuous solution.
- n = The number of target classes which were attacked using noninteger valued strategies in the continuous solution.
- p = The number of weapon types which were allocated using a noninteger valued strategy in the continuous solution
- B_{ki} = The number of type *k* weapons expended by one use of strategy *i*.
- DE_i = The damage expectancy resulting from one use of strategy *i*.
- M = A large positive number.
- S = The set of strategies with noninteger values in the continuous solution.
- ST_j = The subset of *S* that attacks target class *j*.
- SW_k = The subset of *S* that uses weapon type *k*.
- SC = The subset of *S* which attacks target classes that were completely covered in the continuous solution.
- T_j = The number of class *j* targets left to be attacked.
- TC = The set of target classes that were completely covered in the continuous solution.
- W_k = The number of type *k* weapons left to be allocated.
- d^- = The deviation variable associated with the goal.

$$x_i = \begin{cases} 0 \text{ if strategy } i\text{'s truncated value should be unchanged.} \\ 1 \text{ if strategy } i\text{'s truncated value should be increased by 1.} \end{cases}$$

For each target class with uncovered targets after the noninteger valued variables have been truncated, the target constraints (12) ensure that the number of targets attacked is no greater than the number of targets left uncovered. Constraints (13) are the weapon constraints which ensure that, for each weapon type with unused weapons after the noninteger valued variables have been truncated, the number of weapons allocated is no greater than the number of weapons left unused. Constraint (14) is the goal of completely covering those target classes which were completely covered in the continuous solution. The objective function (11) seeks to maximize the total damage expectancy minus a penalty which is accessed for deviating from the goal. The IP constructed has $n + p + 1$ constraints and m binary variables.

The heuristic which uses the Rounding IP to determine which truncated strategies should be increased by one is presented below.

STEP 1: Obtain the continuous solution from AEM. If all variables are integer, stop.

STEP 2: Create a truncated solution by truncating all noninteger valued variables. Set aside this truncated solution.

STEP 3: Gather information from the continuous solution.

a. Identify the target classes which were completely covered in the continuous solution.

b. Identify strategies with noninteger value in the continuous solution.

c. Create a new target list from those targets which were attacked in the continuous solution but are not attacked in the truncated solution. If there is a noninteger number of targets of some class, round up to the next highest integer.

d. Create a new weapons list from those weapons which were used in the continuous solution but are not used in the truncated solution. If there is a noninteger number of weapons of some type, truncate to obtain an integer number of weapons.

STEP 4: Formulate the Rounding IP using information collected at STEP 3.

STEP 5: Solve the Rounding IP.

STEP 6: Combine the solution obtained in STEP 5 with the truncated solution from STEP 2.

HEURISTIC IMPLEMENTATION

In implementing the heuristic, three factors motivated the relaxation of the integer restriction (15) on the decision variables x_i. The first factor is the solution routines used by AEM. AEM uses a linear programming solver and does not contain an algorithm for solving integer programming problems. Switching to another programming code is not practical because of verification and validation concerns and because the current code is tailored to solve the complex, large scale problem of allocating weapons to targets. The second factor is the observation that solutions of the Rounding IP with the integer restriction (15) relaxed often produced integer solutions. The third factor concerns solution times. Since AEM is used by analysts to evaluate a wide range of weapon mixes and targeting constraints, the actual time required to solve realizations of the allocation problem is important, and integer programming problems are often difficult and time consuming to solve. Because of these factors, the integer restriction (15) of the Rounding IP was relaxed. The heuristic implemented in AEM follows.

STEP 1: Obtain the continuous solution from AEM. If all variables are integer, stop.

STEP 2: Create a truncated solution by truncating all noninteger valued variables. Set aside this truncated solution.

STEP 3: Gather information from the continuous solution.

a. Identify the target classes which were completely covered in the continuous solution.

b. Identify strategies with noninteger value in the continuous solution.

c. Create a new target list from those targets which were attacked in the continuous solution but are not attacked in the truncated solution. If there is a noninteger number of targets of some class, round up to the next highest integer.

d. Create a new weapons list from those weapons which were used in the continuous solution but are not used in the truncated solution. If there is a noninteger number of weapons of some type, truncate to obtain an integer number of weapons.

STEP 4: Formulate the LP relaxation of the Rounding IP using information collected at STEP 3.

STEP 5: Solve the LP relaxation of the Rounding IP.

STEP 6: Produce an intermediate integer solution by rounding all noninteger valued variables to the nearest integer.

STEP 7: Combine the intermediate integer solution obtained in STEP 6 with the truncated solution obtained at STEP 2.

At Step 6, the solution of the LP relaxation is rounded. This rounding procedure did not produce infeasible allocations. At Step 7, the truncated integer solution from Step 2 is combined with the integer solution obtained from solving the LP relaxation and rounding at Steps 5 and 6. The combination is accomplished by adding the integer valued strategies from Step 2 to those obtained at Step 6. Once this is done, the new integer solution can be evaluated to determine its damage expectancy, target coverage, and goal performance. A flowchart of the integer solution heuristic is presented in Figure 2.

TESTING

The heuristic has been tested on hundreds of classified real world problems. Experienced analysts report satisfaction with the results. Three unclassified test cases were provided by Mr. William L. Cotsworth and the results of a representative actual case were provided by the US Air Force Studies and Analyses Agency (AFSAA). Information on the three unclassified cases is presented in Table 1. Table 1 shows the case, and for each case, it shows the number of weapons, weapon types, targets, target classes, and goals.

Case 1 strategies could use one or two weapons while cases 2 and 3 allowed only single weapon strategies. Because of the classified nature of the actual case, its information cannot be presented.

For each case, three solutions were produced. In the following discussion, solution 1 is the continuous AEM solution, solution 2 is the integer solution obtained using the current in-

Table 1. The Three Unclassified Cases

Case	Weapons	Weapon Types	Targets	Target Classes	Goals
1	8,691	34	5,774	122	33
2	13,432	19	10,938	11	39
3	29,381	50	34,812	999	11

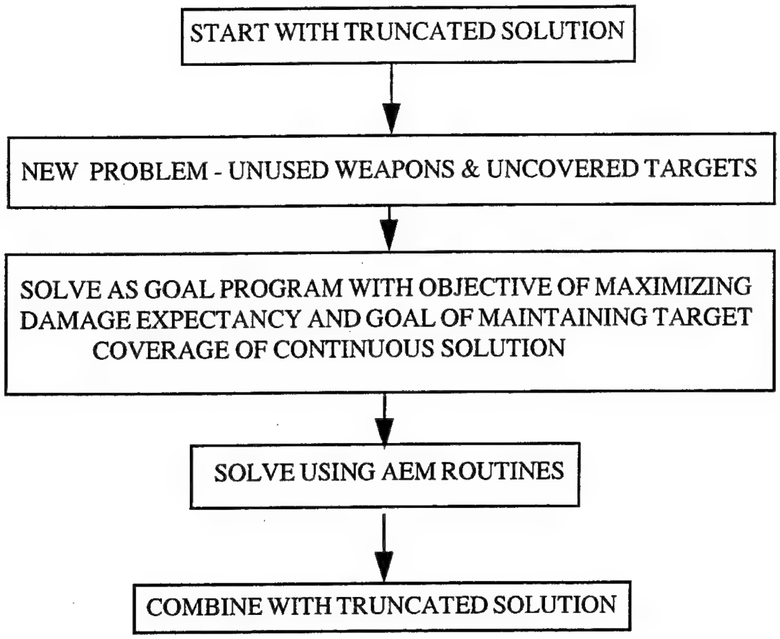

Figure 2. Integer Solution Heuristic

teger solution method in which all noninteger variables in the continuous solution are truncated, and solution 3 is the integer solution obtained using the heuristic presented in the previous section.

Table 2 presents the continuous solution for each of the unclassified cases. The table shows the number of strategies with nonzero value, the number of strategies with noninteger values, the damage expectancy achieved, and the number of targets covered.

Tables 3 and 4 show the damage expectancy and target coverage, respectively, achieved by each of the three solution methods. As can be seen, the damage expectancy achieved by the heuristic (solution 3) matches the damage expectancy obtained by the continuous solution (solution 1) and exceeds that of the truncation method (solution 2). The heuristic covers almost as many targets as the continuous solution and again outperforms the current integer solution approach.

Table 2. The Continuous Solution

Case	Strategies	Noninteger Strategies	Damage Expectancy	Targets Covered
1	160	34	72.59	5503.39
2	98	48	46.46	7225.00
3	1051	46	49.18	19414.23

Table 3. Damage Expectancy

Case	Solution 1	Solution 2	Solution 3
1	72.59	72.25	72.59
2	46.46	46.30	46.46
3	49.18	49.13	49.18

Table 4. Targets Covered

Case	Solution 1	Solution 2	Solution 3
1	5503.39	5481.00	5503.00
2	7225.00	7198.00	7223.00
3	19414.23	19398.00	19413

When the three solution procedures are compared from the view point of goal performance, the heuristic outperforms the current integer solution method and closely approximates the goal performance of the continuous solution. This comparison is presented in Table 5.

The actual case was run at the AFSAA. Because the case is classified, none of the specifics can be given here. Table 6 summarizes the

Table 5. Comparing Results

	Case 1	Case 2	Case 3
Largest Absolute Difference in Goal Performance between Solution 1 and Solution 2	50%	3.75%	0.43%
Largest Absolute Difference in Goal Performance between Solution 1 and Solution 3	0.51%	1.54%	0.04%
Average Absolute Difference in Goal Performance between Solution 1 and Solution 2	3.78%	0.53%	0.11%
Average Absolute Difference in Goal Performance between Solution 1 and Solution 3	0.03%	0.10%	0.01%

Table 6. Performance Summary Actual Case

	$2/1 \times 100$	$3/1 \times 100$
Average Goal Performance	96.18	100.02
Damage Expectancy	99.26	100.01
Target Coverage	99.17	100.01

performance of solutions 2 and 3 as a percentage of the performance of the continuous solution. It can be seen that solution 3 (the heuristic) approximates the continuous solution more closely than solution 2 in terms of goal performance, damage expectancy, and target coverage.

In all reported applications of the heuristic presented above, solution times were extended no more than five seconds. Depending on problem size, AEM can require 10 to 50 minutes of computer time to reach a solution. The five seconds needed to produce superior integer solutions is acceptable to all users of AEM.

CONCLUSIONS AND RECOMMENDATIONS

The integer solution heuristic for the Arsenal Exchange Model developed in this research has been incorporated in AEM. AEM has a user base of 45 Department of Defense organizations and the organizations using the new procedure include United States Strategic Command, J-5 and J-8 of the Joint Staff, and the Air Force Studies and Analyses Agency. These organizations report confidence in and a high level of satisfaction with the integer solution heuristic.

In the four cases tested, the heuristic proved to be an improvement over the current method in terms of damage expectancy, target coverage, and goal performance. Because the heuristic was implemented using the LP relaxation of the Rounding IP, the result is not guaranteed to be integer. In fact, because of the AEM solution procedure, the results will generally not be integer. However, in the cases tested, all the variables were close to integer and, when the noninteger variables were rounded to the nearest integer, the resulting allocation was feasible. If the variables are not all close to integer, it is possible that a weapon or target could be created in the rounding process, so the resulting solution would be infeasible. For example, if two strategies attacking the same target were both selected 0.5 times, and then both rounded up so that they were each selected once, the resulting allocation would attack one more target than was available.

Adding an integer solver to AEM would avoid this problem; however, the risk of producing an infeasible solution may not be great enough to justify the added expense. More research needs to be done to determine under what circumstances an infeasible solution may be produced and the likelihood of these circumstances occurring. Until this problem is solved, if the solution of the LP relaxation is not close to integer, the user should check carefully to ensure that the allocation produced by our approach is feasible.

The new integer solution heuristic we have developed improves the overall performance of AEM. Since AEM is the primary solution engine for CTEM, our procedure can be used to enhance and improve the performance of this model. Because of the general nature of our procedure, it has the potential to be adapted to any model which provides continuous solutions of integer programming problems where resources are assigned or allocated.

ACKNOWLEDGMENTS

We would like to acknowledge the assistance given by the Air Force Studies and Anal-

yses Agency. We would especially like to thank Major Justin Moul for his support and guidance. A special note of thanks goes to Mr. William L. Cotsworth, president of AEM Services Corporation, for his inputs, support, and guidance. We thank the anonymous reviewers whose comments led to an improved paper.

REFERENCES

Benders, J. F. and J. A. E. E. van Nunen. 1983. "A Property of Assignment Type Mixed Integer Linear Programming Problems," *Operations Research Letters*, 2: 47-52.

Bozovich, J. F., W. L. Cotsworth, S. D. Garrett, and P. L. Groover. 1993. The Arsenal Exchange Model (AEM) User's Manual. Englewood, Colorado: AEM Services.

Cotsworth, William L. 1993. President, AEM Services Corporation, Englewood CO. Personal Interview.

Cotsworth, William L. and S. D. Garrett. 1991. The Arsenal Exchange Model Advanced Course. Englewood, Colorado: AEM Services.

Gallagher, Mark A. and Elizabeth J. Kelly. 1991. "A New Methodology for Military Force Structure Analysis," *Operations Research*, 39 (6): 877-885.

Trick, Michael A. 1992. "A Linear Relaxation Heuristic for the Generalized Assignment Problem," *Naval Research Logistics*, 39: 137-151.

Military Modeling for Decision Making, Third Edition

Edited By Wayne P. Hughes, Jr., FS, Captain USN, Retired

There comes a time in every profession when it is appropriate to record the core wisdom generated by and shared among the practitioners and theorists of their craft. After 40 years of military operations research being practiced, MORS felt that this time had come and, in 1989, published the first edition of *Military Modeling*. After selling out of the first and second printruns, MORS decided that, to stay abreast with the computer generation and of the modern modeling tools, a completely revised third edition was in order!

This monograph is the latest step in MORS' dedication to furthering the technical base and professionalism of our membership and the MOR community at large. The purpose of MORS is "to enhance the quality and effectiveness of classified and unclassified military operations research." A key medium for accomplishing that purpose is professional publications such as *Military Modeling for Decision Making*. Through these publications we are enhancing the knowledge of our current members, and more importantly, providing a collated reference for the young analysts yet to join our ranks. It is frankly the latter which motivated the MORS volunteers and staff who labored to produce this monograph. Professor Wayne Hughes, father of this monograph in all its editions, and his group of authors willingly devoted a significant quantity of their scarce time and profound professional knowledge and experience to author and/or revise chapters to reflect the latest technologies, techniques, and applications. We are all quite proud of the result, and trust that it will be used for many years to come by analysts as they employ *Military Modeling for Decision Making*.

Please Send Me _____ copies of Military Modeling for Decision Making at $40 per copy plus $5 for freight.

Name ____

Address ____

____

____

City ______________________ State __________ Zip Code ________

VA Residents add $1.80 sales tax per copy

❑ **CASH** ❑ **CHECK** ❑ **Visa** ❑ **MC** ❑ **AMEX** **Credit card #** ______________________ **Expiration date:** __________ **Billing zip code:** __________ **Name on Credit Card (print):** ______________________ Phone Number: ______________________ **Signature:** ______________________	TOTAL $

101 South Whiting Street • Suite 202 • Alexandria, VA 22304 • (703) 751-7290 • Fax (703) 751-8171
E-mail morsoffice@aol.com

ON THE PRINCIPLE OF "FALSE RANGING" IN WEAPONRY

Israel David

Department of Industrial Engineering and Management, Ben-Gurion University of the Negev, Beer-Sheva, ISRAEL

ABSTRACT

"False ranging" refers to deliberately controlling a gun, or any weapon system, in a counter doctrinal manner, so as to account for its usage in exceptional or improvised missions. Though improvisations and counter doctrinal uses may be debated, and their desirability may change from one army to another, the advantages of "false ranging" are often conspicuous. In this paper we introduce the subject via concrete examples, from some simple ones to the more complex, and then discuss several quantitative issues pertinent to the technique. Devising false ranging menus involves finding simple tables of constants to replace various continuous parameters (unknown to the user in the field). After building an abstract optimization model which is general enough, we demonstrate its use in devising false-ranging menus.

1. INTRODUCTION

Consider the following problem. Army tanks are usually equipped with light mortars, serving as an anti-personnel defensive weapon. Suppose that the mortars are rigidly attached to the turrets of the tanks, and that it has been surprisingly discovered that reasonable terrain obliquity (side inclination) may have significant influence on the impact point of this weapon. There are now several possible attitudes towards the problem: one is to ignore it, arguing that this direct fire may be corrected for by the gunners. A second one is to launch a large-scale emergency project of providing independent balance- mechanisms for all the mortars on the tanks. A third approach is the one which this article is concerned with: it calls for supplying the tank-gunners with a table of constants, specifying what incremental ("false") values they have to add to the true range and to the traversing of the turret so as to account for the side inclination and to hit the target. We call this "false ranging". Note that this approach is not mutually exclusive with the other two. It is only that the problem is not completely ignored, on the one hand, and the orderly solution that should follow need not be an urgent one, on the other hand. A temporary (but satisfactory) solution is easily supplied. Indeed, the calculation of the constants is a matter of simple trigonometry and of examination of the relatively simple flight equations of the short-range mortar bombs. Even more importantly, only very few constants need to be conveyed: all possible combinations of target ranges and side inclinations may be classified to a small number of groups, one constant applying for each group. This is because the accuracy issue is indeed not that critical in the present application.

"False-ranging" type solutions are known in infantry weaponry applications as well as in armor applications, especially for systems with non-computerized fire control. Thus, for instance, tanks with range drums that do not possess an independent scale for *smoke* projectiles are instructed to use an existing scale together with an appropriate "false" range table. Again in tank gunnery, appended "false" range-tables enable firing at *exceptional angles of sight*. (Upwards or downwards direct fire in mountenous areas). Sufficient documentation of such instances in the confines of the present expository paper is problematic either because of classification issues or the length required to appropriately convey the case. Consequently, in Section 2 below we contrive a scenario involved with a *sniper-gun* application, and work out the pertinent false ranging solution in detail. The question is whether it is possible to use an existing sniper-gun system with an improved, newly-introduced, match bullet. Our aim in the presentation is threefold: first to introduce the approach of indirect false-ranging, that is, the use of an auxiliary physical parameter. Second, to show that applications that require high accuracy and are involved with flight equations are soluble in the proposed method, and last, on the presentational level, to equip the reader with a nontrivial case he can follow and work out himself without any need of believing computer simulation reports.

Our last example, and the most complex one, deals again with sudden availability of (improved) ammunition, this time of artillery shells. It shares with the previous example its concern for utmost accuracy (though now for far more distant ranges) and its solution via indirect inputs. The question is how to enable fire of compatible-caliber artillery ammunition which the artillery units *battery computers systems* do not know. This unorderly scenario may arise when civilian companies introduce improved ballistics or extended range ammunition, when ad-hoc coalitions of multinational forces are formed, when a superpower grants a big supply of ammunition to an ally (a decision which is made by politicians and is often unpredicted), when such supply falls to the hands of the army

as capture in wartime, etc. It is to be stressed that even if the full flight simulation modeling of the new projectile is at hand, its introduction to all of the guns battery computer systems in an orderly retrofit process is very slow. Thus, in any of the aforementioned scenarios, and especially in time of emergency, a false ranging solution might be called for. However, in this case a solution in terms of a list which specifies "false" ranges per se is absolutely impractical because of its resultant length. Thus, in analogy with the basic approach of the sniping-gun example, a use of a relevant auxiliary parameter is called for. One could think of introducing to the battery computer a false value of, say, the gun muzzle-velocity parameter, so as to compensate for the "wrong" projectile. An important new characteristic of this problem is that in the case of artillery fire, when aiming to a given target, grossly varied trajectories may be in effect, dependent on the many relevant unknown environmental conditions. Thus, firing at a given gun-elevation and using any given false-ranging parameter would entail variability of the *mean point of impact*. (We shall use throughout the common abbreviation *mpi* for the mean point of impact). This variability should be properly controlled.

The above examples are representative of many others in which false-ranging can be used in exceptional or improvised missions. As a rule, improvisations and counter doctrinal uses should be avoided, and their desirability depends on the spirit of the specific army, on military tradition and on the availability of technological and financial resources. (In addition, the frequency of engaging emergency missions is obviously of major relevance here). At any rate, it is evident that the false ranging principle is versatile and that its use, with appropriate safeguards, may enable a wider and more flexible utilization of a system. In emergency times, it may amount to a decisive contribution in battlefield.

In this paper we discuss several quantitative issues pertinent to devising false ranging in weaponry. In typical applications the false ranging menus are composed of simple tables of constants which are aimed at replacing in the best possible way various continuous parameters (unknown to the user in the field). In the next section we pinpoint this issue by briefly analyzing the small-arms example in a way that the reader could easily check. Since the phenomenon of systematic errors (which in many cases amount to the mpi variability such as in the above artillery case) is common to many weapon systems, it is this issue which will be the focus of the rest of the exposition: in any false ranging method the resultant mpi variability should be controlled for any given entry parameter (say, target range), and the variability over all the relevant entry values be minimized. The optimization and control should be subject however to the requirement of *simplicity*, which is mandatory in manual (non computerized) false ranging. In section 3 the artillery problem will be defined more specifically and some pertinent data, obtained by relevant simulation studies will be cited regarding this problem. In section 4 we build an abstract optimization model to deal with false ranging problems that entail mpi errors in individual entries. Consequently, the use of the model in devising false-ranging menus is discussed in section 5 and is worked out for the data presented in Section 3.

2. THE SNIPER-GUN EXAMPLE

Suppose a certain sniping rifle shoots bullet-type X to a maximum effective (horizontal) range of 1000 meters. The design of the sniping telescope sight is based naturally on the relevant *firing table*, that is, relationship between gun elevations and target ranges. In reference [3] we present a simple explicit model which is accurate enough for the present application and enables one to easily follow the subsequent analysis. We have:

$$\alpha = \frac{1600gx}{3\pi v_0^2(1-kx)}\left(2+\frac{1}{1-kx}\right) \qquad (1)$$

where x is the horizontal target range, α is the elevation (gun firing angle) in mils, g is the coefficient of gravity, v_0 is the bullet muzzle velocity and k is a constant that encapsulates the exterior ballistics of this bullet-type. The physicality of this model is justified in [3] on the grounds of a few simplifying assumptions for flat trajectories, and of empirical findings.

Let us assume that bullet X has v_0=785 meters/sec and $k=k_X=420\times10^{-6}$. (In real life problems this value can be extracted by a least squares fit of the model (1) to firing test results). Columns (1) and (2) of Table 1 below give a partial firing table compatible with (1).

Suppose now that a new, same-caliber match-bullet Y has been introduced, with improved ballistics such that its maximum effective range is 1100 meters (a considerable advantage in sniping rivalry). Suppose further that both bullet types have the same muzzle velocity, and that for the new bullet Eq. (1) holds with $k=k_Y=352\times10^{-6}$. The detailing in Table 1, columns (1) and (3) may now provide a listing of "false ranges" which should be implemented to the sight, so as to hit any range with bullet Y. For example, to hit a target at range 1100 with bullet Y a false range of approximately 1000 should be set. This menu entails however a tedious listing of all ranges.

A different approach might take advantage of fire-control systems which are currently considered by various developers of modern sniper guns. Assume that measured surface air-pressure may be fed to the system so as to be corrected for by the sight. Now, with regard to Eq. (1), further ballistic analysis may show how a nonstandard air pressure may be accounted for compatibly with model (1). We shall use the model:

$$\alpha_x = \frac{3200g}{\pi k^2 xq(q+1)v_0^2}\cdot\left(\frac{1}{(1-kx)^q} - kqx - 1\right) \tag{2}$$

where $q=4c^2-2$, and c is the ratio between the measured air pressure and the standard air pressure, the latter being the value which underlies (1). (For an easily-read ballistics source book the reader is referred to [5]. The heuristic model (2) is advocated in [3]). Again, for the sake of presentation we assume that the fire control system uses a model which agrees with (1), together with the update for air-pressure (2), and that it is familiar with k_X, so to say, but not with k_Y.

Column (4) of Table 1 presents the application of the model (2) with $v_0=785$ m/sec, $k=k_X$ and an air-pressure ("false") value of $c=0.902743$. This value has been obtained by least squares minimization of the respective meter differences from the target, in perpendicular planes at the specified target ranges. These differences are listed in column (5) and are easily calculated from the angular differences (4)-(3) and the range (1). The total optimal sum

Table 1. Accuracy of small-arms false ranging

Range (meters)	α (mils) $k=k_X$	α (mils) $k=k_Y$	α (mils) $k=k_X$ $c=.90274$	dheight (meters) (4) vs. (3)
(1)	(2)	(3)	(4)	(5)
300	2.917	2.827	2.817	−0.004
500	5.586	5.270	5.244	−0.013
700	9.155	8.352	8.309	−0.029
800	11.416	10.208	10.163	−0.036
900	14.108	12.331	12.292	−0.036
1000	17.353	14.777	14.761	−0.017
1100	—	17.619	17.658	0.041

of squares is 55.6 cm^2 and the root-mean-squares (*rms*) is 2.82 cm. Still, taking a step forward, it can be shown that if the two groups of ranges 300 to 800 and 900 to 1100 are taken separately, using two distinct optimal values of c, the achievable accuracy is of 2.12 cm overall rms. (Use $c=0.907291$ for the first group and $c=0.902473$ for the second group). Obviously, the greater the number of c-values employed the more accurate is the method, at the expense of its simplicity however. The final decision as to the exact recipe of false ranging and to the degree of the desirable overall accuracy depends on the specific problem. In the present case relevant considerations may concern the required accuracies of match-fire, the *round to round* precision of bullet types X and Y, and so on. In general we ask what minimal number of c-values must be prescribed, what their values are and what the entailed mpi-accuracies are.

3. AN EXAMPLE WITH MPI VARIABILITY FOR INDIVIDUAL ENTRIES

The purpose of this section is to present an authentic case with *mpi variability*, together with some concrete numerical data. We make specific the conditions of the artillery application that was posed in the introduction, relating to a problem which was solved successfully by false ranging in the past. (Reference [2] contains a detailed unclassified account of that successful application).

Suppose that an artillery force, equipped with modern battery-computer systems, is trained to deliver fire-for-effect of projectile B1 based on prior target-registration with projectile A. This data transfer and non-self registration is supported by the computer and is done in the first place because projectile A is cheap, old-fashioned and less effective, while projectile B1 is modern, effective and expensive. Projectile B2 is a new, even more effective version of projectile B1 which has become available to the military forces in an exceptional procedure such as described in the introduction. A temporary trickery or false input to the existing systems is sought so as to enable immediate use of ammunition B2.

The false ranging solution to the problem has been performed via the aforementioned auxiliary parameter of muzzle velocity. For any specific target and environmental conditions it is possible to find, by numerical solution of the flight equations, the muzzle velocity value that would compensate for the difference between projectile B2 and projectile B1. However, as explained in the introduction and unlike the state of affairs in the simple small-arms example, different muzzle velocity values would fit the same target for different possible environmental conditions, which in this application may include meteorological data, laying and positioning displacements of the guns, and many others. An exhaustive simulation study has therefore been performed to estimate the average and the standard deviation of the resulting muzzle velocity correction (increment to the true muzzle velocity) for individual entry values of the desired false-ranging table. The flight trajectories have been computed using the Modified Point-Mass Trajectory model, see [7] and [8]. For details of the overall simulation algorithm see [2]. Sample results of these simulation studies are given in Table 2, and interpolation can be used to extend these data for any required entry value. (It is to be noted that the natural entry parameter to the false-ranging menu in the present application is the projectile-A final elevation. It is the parameter which becomes known to the operator at the relevant moment and which signifies target registration).

Later on, in Section 5 below, we exemplify our proposed false ranging optimization scheme by applying it to the present data.

Table 2. Optimal muzzle velocity-corrections: numerical results for selected entries.

projectile A elevation (mils)	muzzle velocity-correction		approx. target-range (m)
	average (m/sec)	standard-dev. (m/sec)	
(1)	(2)	(3)	(4)
250	−2.7	3.3	11700
300	−2.7	2.7	13000
350	−4.0	2.4	14000
400	−6.1	2.3	15000
450	−6.6	1.6	15800
500	−7.1	1.7	16600
550	−8.7	1.7	17300
600	−9.1	1.7	17800

4. A GENERAL MODEL FOR FALSE-RANGING WITH MPI VARIABILITY

In this section we propose an abstract model for false ranging that allows for mpi variability in individual entries. This model suggests an appropriate optimization scheme for finding simple and accurate false-ranging tables, as shown later.

We distinguish between the *false-ranging* parameter (e.g. added-range, air pressure and muzzle velocity in the opening examples) and the *entry* parameter. Though we have mentioned a case where the entry parameter could be two dimensional (side-inclination together with target-range in the armour example) we shall assume, for the sake of simplicity, that it is one-dimensional and that it may vary in a given range of values $[a,b]$.

Let us use the following notation:

x - An entry value for the false-ranging table.

Φ - A multidimensional "environmental" variable which is the source of the mpi variability and which will be called "a profile" for short.

$v(x,\varphi)$ - The compensating false-ranging constant, for given x and a realization φ of Φ.

$\mu(x)=E\Phi[v(x,\Phi)]$ - The average optimal false-ranging constant (cf. Table 2).

$s(x)=\sigma_{\Phi}v(x,\Phi)$ - The standard deviation of the constant (cf. Table 2).

$\Delta(u,x,\varphi)$ - The mpi miss which results from using an arbitrary value u of the false-ranging parameter. We have $\Delta(v(x,\varphi),x,\varphi)=0$.

In the last example x is the registration elevation for projectile A, Φ is the meteorological profile, laying and positioning conditions, etc., and $v(x,\varphi)$ is the muzzle-velocity constant which results from the computation which was mentioned in the previous section.

Let us now consider a given range of values of the entry parameter, say x_1 up to x_2.

Definition 1. (Criterion of optimality).

A false-ranging constant u* will be called optimal if:

$$E_{\Phi,\, x_1 \geq X \geq x_2}[\Delta^2(u,X,\Phi)] \tag{3}$$

is minimized by $u=u^*$.

The criterion (3) is a classical weighted least-squares criterion, applied to the target miss using one false-ranging constant for a range of values of the entry parameter. See e.g. [6] for an account of the (deterministic) continuous-parameter least-squares optimization approach. To simplify the exposition, equal weighting for each entry value will be assumed. Thus, the expression in (3) becomes:

$$\int_{x=x_1}^{x_2} \frac{1}{x_2 - x_1} E_{\Phi}[\Delta^2(u,x,\Phi)]dx. \tag{4}$$

We now proceed to find formulae for the optimal false-ranging-constant and for the accuracy of its application in a range of entry values.

For the expectation in (4) we have:

$$E_{\Phi}[\Delta^2(u,x,\Phi)] = Var_{\Phi}[\Delta(u,x,\Phi)] + E^2_{\Phi}[\Delta(u,x,\Phi)]. \tag{5}$$

Writing now u_0 for $v(x,\phi)$ and neglecting nonlinear terms of the relevant Taylor expansion we get:

$$\Delta(u,x,\varphi) = \Delta(u_0,x,\varphi) + (u - u_0)\frac{\partial}{\partial u}\Delta(u,x,\varphi)|_{u=u_0}. \tag{6}$$

The first term in the R.H.S of (6) equals zero.

We stipulate that for any x, a kind of a "mean value" profile φ_x exists, in a way that the following hold:

$$\mu(x) = v(x,\varphi_x). \tag{7}$$

In fact, in many applications a "profile" includes many factors, such that there might be many solutions φ_x for Eq. (7). We emphasize, however, that φ_x plays only a conceptual role in the subsequent analysis and its determination is not required in practice. Defining now:

$$D(x) = \frac{\partial}{\partial u}\Delta(u,x,\varphi_x)|_{u=u_0} \tag{8}$$

we get:

$$\Delta(u,x,\varphi_x) = (u - \mu(x))D(x). \tag{9}$$

A few words are in order here. $D(x)$ is the incremental miss of the target incurred by an incremental change in the optimal false-ranging-parameter. This function of x has to be assessed by the analyzer. In the last artillery problem $D(x)=p_1(x)\cdot m_1(x)/p_2(x)$, where $p_1(x)$ and $p_2(x)$ are the incremental range per incremental gun elevation for projectile B1 and projectile B2 respectively, and $m_1(x)$ is the incremental range per incremental muzzle-velocity for projectile B1. All three of these functions are available from the respective standard books of firing tables. (For the contents and the standard NATO format tables see e.g. [8], or any US Army book of field artillery firing tables).

In order to get an applicable expression for (5) we proceed in making three assumptions. First,

$$E_{\Phi}[\Delta(u,x,\Phi)] = \Delta(u,x,\varphi_x), \tag{10}$$

that is, the mean mpi-error is attained at the mean profile (see remark after (7) above). Second,

$$\sigma_{\Phi}[\Delta(u,x,\Phi)] = \sigma_{\Phi}[\Delta(\mu(x),x,\Phi)] \tag{11}$$

(The *variance* of the mpi-error while implementing a given false-ranging value u is the same as in implementing the optimal u. This, unlike the *mean* error, is approximately true for values of u which are not too-distant from $\mu(x)$). Finally,

$$\sigma_{\Phi}[\Delta(\mu(x),x,\Phi)] = \Delta(\mu(x) + s(x),x,\varphi_x), \tag{12}$$

namely, the standard error using the optimal false-ranging constant approximately equals the error (miss of target) which is incurred (for an average profile) when using a false-ranging value which equals the average optimal constant plus its standard deviation. Combining (5)

and (9)-(12) we have:

$$E_{\Phi}[\Delta^2(u,x,\Phi)] = D^2(x)(s^2(x) + (u - \mu(x))^2) \tag{13}$$

By Definition 1 we get the optimal u by differentiating (4) and equating to zero. Using (13) and Leibnitz' rule:

$$\frac{d}{du}[\text{expression in}$$

$$= \frac{1}{x_2 - x_1}\int_{x_1}^{x_2} \frac{d}{du}[D^2(x)(s^2(x)$$

$$+ (u - \mu(x))^2)]dx = \frac{2}{x_2 - x_1}$$

$$\cdot \int_{x_1}^{x_2} (u - \mu(x))\, D^2(x)dx$$

and finally

$$u^* = \frac{\int_{x_1}^{x_2} \mu(x))\, D^2(x)dx}{\int_{x_1}^{x_2} D^2(x)dx}. \tag{14}$$

This is the formula for the *optimal constant in a range* $[x_1,x_2]$. We see that it depends on $\mu(x)$ but not on $s(x)$. From (13) we also get the *root-mean-square* accuracy of using u^* for a given x in $[x_1,x_2]$, namely,

$$rms(x) = D(x)\sqrt{s^2(x) + (u^* - \mu(x))^2}. \tag{15}$$

5. DEVISING FALSE-RANGING TABLES IN PRACTICE

Our objective now is to determine a *partition* of the interval [a,b] of the relevant values of the entry parameter. Regarding a partition $a=x_0<x_1<x_2<\ldots<x_n=b$ we say that there are n *zones* (subintervals) and that the x's are *zone delimiters*. For each zone a constant false-ranging-value is to be assigned. We require that the following three conditions will be met:

a - The root-mean-square error will be bounded by a predetermined constant, uniformly for each entry value.

b - The number of zones n will be minimal such that condition a holds, and

c - The total rms error over [a,b] will be the minimum possible for n.

In [4] we present a special-purpose *dynamic programming* algorithm for the exact solution of false-ranging like problems. It is based on discretization of the entry-parameter domain, and on shortest-path finding of the optimal partition for given number of subintervals n. (Optimality in the sense of least total rms). A similar procedure has already been suggested by Bellman and Dreyfus [1, pp. 334-335] for the basic problem of the optimal least-squares approximation to a function by a piecewise-linear function. They assumed a given number n of linear segments. Though the present interpretation and the overall least-squares functional are different, the logic of the dynamic programming solution is identical to that of [1]. Starting with $n=1$, the resultant optimal solution is checked for *feasibility* (condition a above). If this check fails, n is increased by one. The procedure stops when a feasible (condition a) - minimal (condition c) solution is found for the first time. Minimality in the sense of condition c is guaranteed by exploiting Eq. (14) within the dynamic programming algorithm. (This last fact implies a need for an extra check of the resultant partition, namely a check for the validity of the assumptions that underlie (14) and are listed in the previous section).

As it happens frequently in military OR, field applications involve additional factors which are not included in the neat model. In this case, for example, we have neglected the need for a standard format for *various* tables (there are several propellant types and charge zones), the desirability of *round* partition delimiters, etc. Attempting a flexible approach, we note that by the nature of the application only a small number of zones have to be determined. Hence it would be practical to find the required partition and the optimal u^*'s (the "menu") by *trial and error*. This could be done using the handy formulae of the previous section. The results of Section 4 show how to find the rms-optimal false-ranging value u^* for any subinterval $[x_i,x_{i+1}]$, and what is the incurred rms error for any entry value x when using u^*. Furthermore, for validating the boundedness condition a we content ourselves with checking rms(x)

only at the zone delimiters. Indeed, recognizing u^* (14) as a *mean value* of $\mu(x)$ on the respective zone, it is plausible that more often than not rms(x) would be bounded from above by the maximum of rms(x_1) and rms(x_2), for all x in $[x_1,x_2]$. For example, if $D(x)$ is constant, then by (14) u^* is a mean value strictly by the *mean-value theorem*. If, in addition, $\mu(x)$ is linear and $s(x)$ is constant we get under the root in (15) a parabola which is minimized at the mean value, and maximized at the end values x_1 and x_2.

The necessary inputs to carry out the optimization are recognized as the functions $\mu(x)$, $s(x)$, and $D(x)$ of Section 4. All of these functions are supposed to be continuous on [a,b]. We remark parenthetically that the solution for such cases as discussed in Section 2 come out as a special case of the present scheme, with $s(x)=0$. Note that for the small-arms example $\mu(x)$ equals $c(x)$, and $D(x)$ can be analitically derived from Eq. (2).

Let us work out the case of Section 3. $\mu(x)$ and $s(x)$ are interpolated from the values of Table 2 for $250 \leq x \leq 600$, where x is projectile-A elevation (in mils). The function $D(x)$ is evaluated as explained in Section 3 (relevant values appear in Table 3 below). Suppose that an mpi variability bound of 50m rms is required. By inspection of Table 2 one may converge to the following table after a short trial and error process:

In Table 3 column (1) contains the tentative partition, and column (2) states the zone-optimal false-ranging constants according to Eq. (14). The next three columns show the necessary inputs that were discussed above, and the last column states the zone-delimiter rms according to Eq. (15). Although the method of trial and error does not necessarily minimize the total rms error, we obtain a convenient partition which obeys condition a. It can be seen that no significant deviation from the proposed partition is possible, so that the optimal number of zones (and of false-ranging constants) must be 3.

Even if overall accuracy can be somewhat improved it is advantageous to have round delimiter values, and therefore the present table is satisfactory. In fact, applying the dynamic programming algorithm [4] for the current ex-

Table 3. Optimal u's and delimiter rms for a tentative false-ranginge table

projectile-A elevation zone delimiters (mils)	zone optimal u^* (m/sec)	μ (x) (m/sec)	s (x) (m/sec)	D (x) (m/sec)	rms (x) (m)
(1)	(2)	(3)	(4)	(5)	(6)
250	−2.7	−2.7	3.3	14.5	47.9
300		−2.7	2.7	15.2	41.0
300	−4.3	−2.7	2.7	15.2	47.4
400		−6.1	2.3	16.9	49.5
400	−7.6	−6.1	2.3	16.9	46.5
600		−9.1	1.7	21.4	47.9

Table 4. Optimal elevation zones, optimal u's and delimiter rms

projectile-A elevation zone delimiters (mils)	zone optimal u^* (m/sec)	μ (x) (m/sec)	s (x) (m/sec)	D (x) (m/sec)	rms (x) (m)
(1)	(2)	(3)	(4)	(5)	(6)
250	−3.3	−2.7	3.3	14.5	48.6
370		−2.7	2.7	15.2	46.0
370	−6.4	−2.7	2.7	15.2	46.6
510		−6.1	2.3	16.9	36.9
510	−8.6	−6.1	2.3	16.9	38.6
600		−9.1	1.7	21.4	38.2

ample yields the following optimal solution, up to a resolution of 10 mils in the entry-parameter:

The overall rms of Table 3 is 40.4 meters, while the optimal rms (Table 4) is 38 meters.

CONCLUSION

It has been argued in this work that the general approach of false ranging may be manifestly rewarding when no orderly and fast solution is possible. While each instance may require its own "trick" (and sometimes real innovative thought), there are clear unifying characteristics of the application of this principle. In general, the design of false ranging involves the following three interrelated questions. Firstly, what are the physically accessible inputs or settings in the system that are candidates for false positioning? Secondly, how can one provide the operator of the system (in cases where the improvisation is not automated) with the shortest and simplest possible ad-hoc menu? Thirdly, what are the performances, usually accuracies, of the improvised method?

In the paper different evolutionary stages of the principle have been mentioned. The most simplistic one, that is, a detailed list of "false" ranges versus the "true" ones is mentioned (in the description of Table 1). The next step, which includes aggregation of values to make a simpler list, comprises the solution for the tank-mortar example which was mentioned in the Introduction. Then the idea of using an auxiliary parameter is introduced (the sniper example) - though still in a deterministic context. Finally, the "noisy", probabilistic application is discussed - applied to the artillery problem.

As was demonstrated in the paper, false ranging recipes typically involve substitution of discrete values for a certain continuously changing parameter. We have adopted a least-squares optimization approach to build an abstract model of enough generality to cope with false-ranging problems with mpi variabilities. Note that the present formulation particularly suits analyses where a perfect analytical representation of the system is unavailable, and yet the average and the standard deviation of the optimal false-ranging value for any single entry value may be assessed by simulation studies. This is often the case when dealing with systems which are technically and operationally complex. The model we built serves in a convenient way in devising false ranging recipes that comply with the simultaneous requirements of accuracy and simplicity, when such recipes are feasible.

ACKNOWLEDGMENT

I wish to thank the anonymous referees for valuable help through several revisions.

REFERENCES

Bellman, R.E. and S.E. Dreyfus, *Applied Dynamic Programming* (1962), Princeton University Press, Princeton, N.J.

David, I., "The Principle of False Ranging in Weaponry and one Problem in Artillery Operations Research", Technical Report, Department of IE and Management, Ben-Gurion University, Beer-Sheva, Israel (July 1995).

David, I., "Approximate Explicit Solutions for Flat Ballistic Trajectories", Technical Report, Department of IE and Management, Ben-Gurion University, Beer-Sheva, Israel (November 1996).

David, I., "A Dynamic Programming Approach to False-ranging type Optimization Problems", Technical Report, Department of IE and Management, Ben-Gurion University, Beer-Sheva, Israel (November 1996).

Farrar, C.L. and D.W. Leeming, *Military Ballistics – A Basic Manual* (1983), Brassey's Publications, Oxford, England.

Isaacson, E. and H.B. Keller, *Analysis of Numerical Methods* (1966), Wiley, New York, N.Y.

Lieske, R.F. and M.L. Reiter, "Equations of Motion for a Modified Point-mass Trajectory", BRL Report no. 1314 (AD 485869) (1966).

Textbook of Ballistics and Gunnery, Vol. 2, Part III (Measurements, Instrumentation and Data Derivation), London, Her Majesty's Stationery Office (1984).

RiskNav™: A Decision Aid for Prioritizing, Displaying, and Tracking Program Risk

C. C. Cho, P. R. Garvey, and R. J. Giallombardo

The MITRE Corporation
Economic and Decision Analysis Center
202 Burlington Road
Bedford, Massachusetts 01730

ABSTRACT

This paper presents a management tool for prioritizing, displaying, and tracking program risk. The tool, called RiskNav, provides program offices a structure for conducting continuous risk assessments. Using classical utility theory, RiskNav rank-orders and tracks project-defined risk events as a function of their estimated cost, schedule, and technical performance impacts. This includes quantifying the effects of coupled risk events. As a decision-aid, RiskNav assists in identifying where engineering resources are best allocated to mitigate potentially crippling areas of risk to a program.

A. INTRODUCTION

Successfully managing today's major defense systems requires deliberate and continuous attention to risk due to reduced budgets, expanded mission objectives, and increased schedule pressures. To support the risk assessment process, a methodology and software application named RiskNav was developed.

RiskNav is a management tool that facilitates frequent program risk assessments. RiskNav follows the risk assessment process in AFMCPAM 63-101, the Air Force's Acquisition Risk Management Guide (AFMCPAM, 1993). As such, RiskNav offers program managers a means to integrate risk assessment directly into their systems engineering processes.

Among the outputs produced by RiskNav are its graphical displays. A primary display is the relative ranking of risk events prioritized by their impact to a program and probability of occurrence. An example is shown in Figure 1.

Additional displays are available that show the level of coupling (dependence) between any two risk events. This identifies where attacking one event may offer the added value of attacking associated events.

The following presents the underlying theory in RiskNav. This includes a method for quantifying the effects of coupled risk events, an issue not well addressed in current practice.

B. METHODOLOGY OVERVIEW

B.1 Process

Implementing RiskNav begins with the formation of a cross-functional project risk assessment team. Membership includes representatives from the major engineering and program control disciplines. Once in place, the team proceeds through the six steps illustrated in Figure 2.

These steps are described as follows:

1. Identify ... Identify key program risk areas.
2. Define ... Define the set of risk events that fall within the risk areas identified in step 1 (this includes defining event coupling relationships).
3. Assess ... Assess subjective probabilities that each risk event will occur.
4. Estimate ... Estimate impacts for each risk event.
5. Display ... Display the prioritization of risk events determined by RiskNav's ranking algorithms.
6. Analyze ... Evaluate results, check for consistency, and conduct sensitivity analyses. Results serve as inputs to the formulation of risk mitigation strategies by the management team.

Risk areas are defined at the functional area of a program. An example might be *Message Processing*. For each risk area defined by the risk assessment team, the set of events considered to have an unfavorable outcome are then specified. These are defined as risk events. For instance, a risk event within the risk area *Message Processing* might be

> "Current processor fails to meet search and beacon plot message processing requirements of 20,000/12 seconds."

Ideally, risk events should be stated with sufficient clarity to support an assessment of the subjective probability of the event occurring. Additional examples of risk events are provided in Section B.5.

A variety of assessment methods can be used for the first four steps. These include the use of independent surveys, one-on-one interviews across the project, and electronic meeting room technologies. Documentation of the risk assessment team's rationale, for each step shown in Figure 2, is a critical, history-preserving, aspect of the process.

B.2 Risk Event Prioritization

RiskNav defines risk events as vectors in n-dimensional space. The components of these vectors represent areas of a program impacted by the event.

APPLICATION AREA:
Resources and Readiness

OR METHODOLOGY:
Probabilistic Operations Research

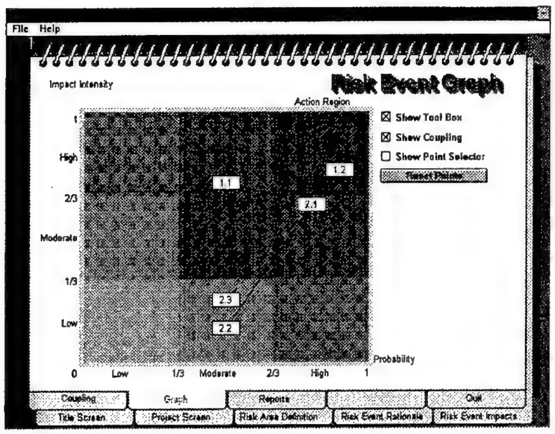

Figure 1. A RiskNav Display.

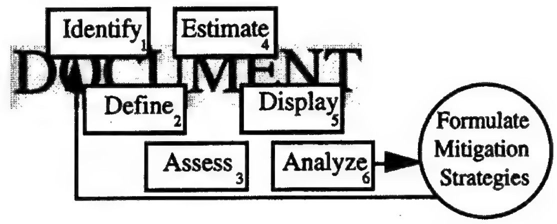

Figure 2. RiskNav Process (Steps 1–6).

Three areas typically impacted by risk events are shown in Figure 3. The components x_C, x_S, and x_T represent the impact of risk event X on a program's cost, schedule, and technical performance areas. To provide a basis for comparison, these impacts are normalized to unit scales.

For convenience, this paper presents RiskNav's methodology in terms of the above three components; higher dimensional generalizations to more than three components* is straightforward.

Risk events are prioritized by their "impact intensity". One way to measure intensity is by the norm of X. Specifically, the impact intensity I of risk event X is

$$0 \le I(X) = \sqrt{(x_C^2 + x_S^2 + x_T^2)/3} \le 1 \qquad (1)$$

where $0 \le x_C, x_S, x_T \le 1$ and x_C, x_S, and x_T are the component-level cost-schedule-technical (CST) impact intensities of X.

Equation 1 is a special case of the Hardy-Littlewood-Polya weighted mean (Hardy, 1967), which, when normalized to the unit-interval, is given by equation 2.

$$0 \le I_A(X) = \left(\frac{1}{W}\sum_{r=1}^{n} w_r x_r^k\right)^{1/k} \le 1 \qquad (2)$$

* In practice, a particular risk event impacts any number of areas (components) of a program. The number and nature of these components must be defined and specified by the risk assessment team.

In the expression above, $k > 0$, $0 \le x_r \le 1$, and w_r is a positive weight assigned to the rth-component x_r. The sum of the component-level weights is denoted by W.

Equation 2 allows the risk assessment team to directly weigh component-level impact intensities. For instance, if a risk event's impact intensity is measured by the rule

$$I_A(X) = \left(\frac{x_C^2 + x_S^2 + 4x_T^2}{6}\right)^{1/2}$$

where $x_1 = x_C$, $x_2 = x_S$, and $x_3 = x_T$, then the event's technical impact x_T is given twice the weight of x_C and x_S. When $k = 1$, the additive impact intensity measure becomes a linear weighted mean of the components. A fixed increment in each component produces a fixed increment in the overall impact intensity, independent of the values of the components. With $k \ne 1$ this is no longer true. The effect on the overall impact intensity from each increment in a component will depend on the value of the component as well as on the values of the other components.

From utility theory, equation 2 (with $k = 1$) reduces to the well-known additive multi-at-

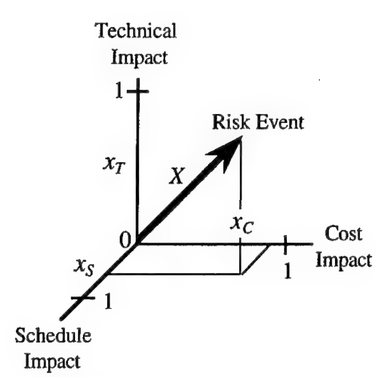

Figure 3. A Risk Event Vector.

tribute utility function (MUF) (Keeney, 1976). However, with equation 2 a high utility (impact intensity) in one component (e.g., x_C) does not necessarily imply a high utility (impact intensity) overall. A component of X at its maximum impact (e.g., $x_C = 1$) does not guarantee $I_A(X)$ will be maximum.

A measurement rule that guarantees the overall impact intensity is maximum whenever a component of X is at its maximum intensity is given by equation 3.

$$0 \leq I_M(X) = [1 - \prod_{r=1}^{n} (1 - x_r^k)^{w_r}]^{1/k} \leq 1 \qquad (3)$$

Equation 3 is related to the well-known Keeney-Raiffa multiplicative MUF (Keeney, 1976). In particular, if $k = 1$, $w_r = 1$ (for all r), $n = 3$, and $x_1 = x_C$, $x_2 = x_S$, and $x_3 = x_T$, equation 3 reduces to

$$I_M(X) = 1 - (1 - x_C)(1 - x_S)(1 - x_T) \qquad (3a)$$

Equation 3a exhibits a property sometimes referred to as constructive interaction. That is, a high utility (impact intensity) in one component (e.g., x_C) means a high utility (impact intensity) overall.

Notice if *any* component x_C, x_S, or x_T (in equation 3 or 3a) is unity, then $I_M(X) = 1$. In equation 3, this is true regardless of the weight assigned to the component. Such a property is desirable for military decision-makers. It allows a risk event to be signaled (flagged) for further consideration when just one component is at its extreme. A comparison of $I_A(X)$ and $I_M(X)$ is given in Table 1.

Measuring a risk event's impact intensity by equations 2 or 3, permits the risk assessment team to define weights as a constant or as function of the component-level intensities. For instance, suppose an additive rule (equation 2) was used to measure impact intensity. The team might define w_1 as an increasing function of x_C as x_C approaches one. Three possible functions are illustrated in Figure 4.

Viewing risk events as vectors in n-dimensional space, whose components reflect the areas of a program impacted by the events, is a useful way to conceptualize risk and plan for its elimination. Measuring an event's impact by

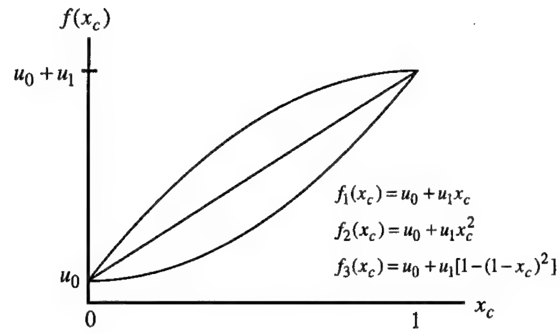

Figure 4. Illustrative Weight Functions for Cost Impact Intensity x_C.

the rules described, provides the flexibility needed by a risk assessment team. Measurement rules, such as equations 2 or 3, can be defined to reflect the degree of risk averseness appropriate to the program and the decision-maker's preference structure.

B.3 Impact Scales

Illustrated in Figure 2, a risk event is viewed as a vector whose components reflect the event's CST impacts to the program. These components are normalized to the unit interval, so that the impact intensities of each event can be compared on a common scale.

Estimating CST impacts is a critical aspect of the risk assessment process. It can also be the most time and resource intensive aspect of the process. The resources applied to estimate these impacts should reflect the purpose for, and scope of, the risk assessment. Clearly, more care

Table 1. Measuring Impact Intensity

Values For (x_C, x_S, x_T)	I_A(X) (Eqt. 1)	I_M(X) (Eqt. 3a)
(0.00, 0.00, 0.00)	0.00	0.00
(0.00, 0.00, 0.25)	0.14	0.25
(0.00, 0.00, 0.50)	0.29	0.50
(0.00, 0.50, 0.50)	0.41	0.75
(0.50, 0.50, 0.50)	0.50	0.88
(0.00, 0.00, 1.00)	0.58	1.00
(0.00, 0.50, 1.00)	0.65	1.00
(0.00, 1.00, 1.00)	0.82	1.00
(0.50, 1.00, 1.00)	0.87	1.00
(1.00, 1.00, 1.00)	1.00	1.00

given to developing impact estimates offers a decision-maker a more meaningful risk assessment.

Unfortunately, it is a frequent case that risk assessments are done with very little time and resources available to estimate CST impacts directly. To facilitate the risk assessment process in such circumstances, a set of utility scales can be developed to "bound" the impacts. Figure 5 illustrates a utility scale developed for technical performance.

This scale was modified from AFMCPAM 63-101 (AFMCPAM, 1993). Further illustrations of utility scales, including those for cost and schedule, can be found in reference 1. Utility scales must, in general, be customized for each specific project application.

B.4 Risk Event Coupling

An important consideration in any risk assessment is the identification of event relationships. Risk events can relate to each other in their occurrences. They may influence each other to occur together (concurrently or sequentially) or in each other's absence, or their occurrences may be totally independent of each other. This relation of association or disassociation, if any, may be traceable to some root causes that make events appear to induce or exclude each other's occurrence. For RiskNav a

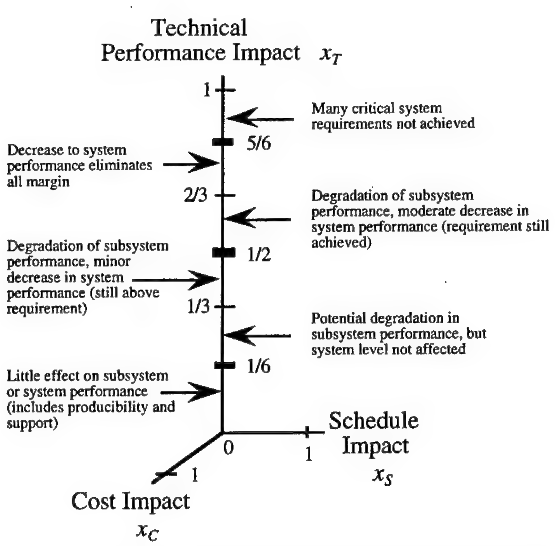

Figure 5. A Utility Scale for Technical Performance Impact.

procedure was developed to calculate "coupling level" which was designed to measure such a relation.

RiskNav defines event coupling as the tendency of the relation between two risk events. A coupling level is defined which measures the strength of the tendency. Coupling level ranges in value between +1 and −1. A coupling level of 0 reflects the case where two events are independent of each other in their occurrences. A positive coupling level between 0 and +1 results when two events have a tendency to *"induce"* each other's occurrence. A negative coupling level between 0 and −1 results when two events have a tendency to *"exclude"* each other's occurrence. Either positive or negative, the two events are dependent on each other. The coupling level defined here is used to measure the strength of their tendency to *"induce"* or *"exclude"* each other's occurrence.

The following provides definitions used in the rest of this paper. In each, it is useful to consider the occurrence of a risk event in a common time period (e.g., between time 0 and time *T*, inclusively) that is relevant to the situation under analysis. If an event occurs after time *T*, it is considered a non-occurrence for the purposes of the analysis. The terms are:

A) $P(X_i)$ is the non-zero marginal probability that risk event X_i will occur. It expresses the likelihood of the event's occurrence at some time t during a period of interest.

B) $P(X_i|X_j)$ is the conditional probability that risk event X_i will occur given that risk event X_j will occur. This does not imply that event X_j already occurred or that it will occur before event X_i. It expresses the likelihood of the occurrence of event X_i if event X_j is known to occur at some time t during a period of interest.

C) $P(X_i, X_j)$ is the joint probability that risk events X_i and X_j will both occur. This does not imply any sequence of event occurrences. It expresses the likelihood that both risk events will occur, each at some time t during a period of interest.
A basic relationship between the above probabilities is given by equations 4 and 4a.

$$P(X_i, X_j) = P(X_i|X_j)\,P(X_j) \qquad (4)$$

$$= P(X_j|X_i)\,P(X_i) \qquad (4a)$$

D) $L(X_i, X_j)$ is the coupling level between the occurrences of risk events X_i and X_j. It

is a measure of the strength of the tendency between the two events' occurrences, each at some time t during a period of interest.

B.4.1 Measuring Risk Event Coupling. The coupling level of two risk events X_i and X_j is measured as follows:
Rule A: If $P(X_i, X_j) > P(X_i)P(X_j)$ then

$$L(X_i, X_j) = \frac{P(X_i, X_j) - P(X_i)P(X_j)}{\text{Min}\{P(X_i), P(X_j)\} - P(X_i)P(X_j)}$$

Rule B: If $P(X_i, X_j) \leq P(X_i)P(X_j)$ then

$$L(X_i, X_j) = \frac{P(X_i, X_j) - P(X_i)P(X_j)}{P(X_i)P(X_j)}$$

The above measurement rules normalize the joint probability $P(X_i, X_j)$ of risk events X_i and X_j by a value that is a function of their marginal probabilities $P(X_i)$ and $P(X_j)$. The normalization reflects that the tendency of the relation between two events' occurrences can be strong even though their marginal probabilities may be low and that the tendency can be weak when the marginal probabilities may be high. A further discussion of this is provided in the appendix. The following conditions result from the coupling measurement rules defined above.

1. If the occurrence of one risk event, say X_j, definitely *"induces"* the occurrence of the other event X_i, that is

$$P(X_i|X_j) = 1 \Rightarrow P(X_i, X_j) = P(X_j)$$

and

$$P(X_j) \leq P(X_i)$$

then

$$L(X_i, X_j) = +1$$

Thus, the coupling level +1 indicates definite *"induction"* between events X_i and X_j. The coupling level is +1 if $P(X_i|X_j) = 1$, or $P(X_j|X_i) = 1$, or both.

2. If one risk event, say X_j, has no influence on the other event X_i (the two events are independent), that is

$$P(X_i|X_j) = P(X_i) \Rightarrow P(X_i, X_j) = P(X_i)P(X_j)$$

then

$$L(X_i, X_j) = 0$$

Thus, the coupling level 0 indicates independence between X_i and X_j. The independence is always mutual since $P(X_i|X_j) = P(X_i)$ implies $P(X_j|X_i) = P(X_j)$.

3. If the occurrence of one risk event, say X_j, definitely *"excludes"* the other event X_i, that is

$$P(X_i|X_j) = 0 \Rightarrow P(X_i, X_j) = 0$$

then

$$L(X_i, X_j) = -1$$

Thus, the coupling level of -1 indicates definite *"exclusion"* between events X_i and X_j. The definite *"exclusion"* is also always mutual since the conditional probability $P(X_i|X_j) = 0$ implies $P(X_j|X_i) = 0$.

To compute $L(X_i, X_j)$ the necessary inputs are the joint probability $P(X_i, X_j)$ and the marginal probabilities $P(X_i)$ and $P(X_j)$.

Values for the marginal probabilities most often reflect subjective assessments. They are made early in the risk assessment process. The joint probability is the remaining input needed to complete the calculation of $L(X_i, X_j)$. Unlike the marginal probabilities, the joint probability is not assessed directly. It is determined from equation 4, which requires an assessment of the conditional probability $P(X_i|X_j)$. There is a range of feasible values from which to make this assessment. This range is a function of the marginal probabilities $P(X_i)$ and $P(X_j)$. It is given by

$$l_{i|j} \leq P(X_i|X_j) \leq u_{i|j}$$

where

$$l_{i|j} = \text{Max}\left\{0, 1 - \frac{1 - P(X_i)}{P(X_j)}\right\} \tag{5}$$

$$u_{i|j} = \frac{\text{Min}\{P(X_i), P(X_j)\}}{P(X_j)} \tag{6}$$

B.4.2 Assessing Probabilities

This section offers guidelines for assessing the marginal and conditional probabilities needed to measure the coupling level between

two risk events. There is a large corpus of literature on the subject of subjective probability elicitation. An excellent treatment of this topic, along with a detailed bibliography, is given by Clemen (Clemen, 1990).

The guidelines below provide a framework for translating qualitative assessments of an event's likelihood of occurrence into probability values. A translation scheme for assigning marginal probabilities is offered in Table 2. The author's experience is the basis for the scheme presented.

When risk event coupling is considered, it is necessary to assess the conditional probability of an event's occurrence given the occurrence of another (related) event. This assessment differs from the marginal probability assessment in two respects. First, the conditional probability assessment is easier if done in relation to the marginal probability, which is assessed earlier in the RiskNav process. Second, there are bounds ($l_{i|j}$ and $u_{i|j}$) that are applicable to the conditional probability. A scheme, also based on the author's experience, is suggested in Table 3. This aids the risk assessment team in making conditional probability assessments as they relate to risk event coupling.

B.5 An Illustration

This section illustrates portions of a risk assessment conducted on a radar system modernization. Table 4 presents several RiskNav inputs and outputs from that assessment.

Table 2. A Marginal Probability Assignment Scheme

Qualitative Assessment of Event X_i Occurring	Quantitative Translation $P(X_i)$
Sure to Occur	1.0
Almost Sure to Occur	0.9
Very Likely to Occur	0.8
Likely to Occur	0.7
Somewhat > than an Even Chance	0.6
An Even Chance to Occur	0.5
Somewhat < than an Even Chance	0.4
Not Very Likely to Occur	0.3
Not Likely to Occur	0.2
Almost Sure Not to Occur	0.1
Sure Not to Occur	0.0

Table 3. A Conditional Probability Assignment Scheme

Qualitative Assessment of Event X_i Occurring Given the Occurrence of X_j	Quantitative Translation $P(X_i\|X_j)$
Likelihood Increased to Maximum Possible	$u_{i\|j}$
Likelihood Increased Significantly	$P(X_i) + \frac{u_{i\|j} - P(X_i)}{2}$
Likelihood Basically Unchanged	$P(X_i)$
Likelihood Decreased Significantly	$P(X_i) - \frac{P(X_i) - l_{i\|j}}{2}$
Likelihood Decreased to Minimum Possible	$l_{i\|j}$

Table 4 summarizes the key inputs required for RiskNav. These inputs must be developed by the risk assessment team. They include establishing the impact scale maximums, estimating CST impacts (x_C, x_S, and x_T) for each risk event, and assessing the subjective probabilities of each risk event (i.e., the marginal probabilities).

Table 4 also presents the impact intensities computed for each risk event. In this example, equation 1 was used. The risk events having the highest impact intensities are $X_{1.1}$, $X_{1.2}$, and $X_{2.1}$. A display generated by RiskNav, of these intensities is provided at Figure 1. This shows the relative ranking of these events prioritized by their combined CST impacts and probabilities of occurrence.

Next, it is important to identify risk event interrelationships. For purposes of this illustration, suppose the risk assessment team identified that coupling relationships exist between the four pairs of risk events shown in Table 5. From the marginal probabilities assessed for each risk event in Table 4, coupling levels $L(X_i, X_j)$ are determined from either rule A or rule B.

Figure 6 presents a RiskNav display designed to illustrate risk event coupling and the computed coupling levels. Such a display provides decision-makers an important visual aid; one that focuses their attention on attacking critical risk events that also have high-impact associations with other coupled events under consideration. Figure 6 illustrates a moderate to high tendency for risk event $X_{2.1}$, if it occurs, to

Table 4. Radar System Modernization An Example RiskNav Risk Assessment

Impact Scale Maximums*	$\text{Max}\{x_C\} = 1$ at \$20M $\text{Max}\{x_S\} = 1$ at 36 Mos. $\text{Max}\{x_T\} = 1$ (by defn.)
Risk Area 1. This area includes two risk events associated with the system's ability to meet surveillance capacity requirements.	Risk Event $X_{1.1}$. Message Processing—"Currently proposed processor fails to meet search and beacon plot message processing requirements of 20,000/12 seconds." Risk Assessment Team Inputs: $x_C = \frac{17}{20}, x_S = \frac{15}{36}, x_T = 0.8, P(X_{1.1}) = 0.5$ RiskNav Output: *Impact Intensity I = 0.72*
Risk Area 1 (Concluded)	Risk Event $X_{1.2}$. Track Processing—"Currently available COTS/GOTS fails to meet requirement for processing/fusing internal and external tracks into a single track." Risk Assessment Team Inputs: $x_C = \frac{16}{20}, x_S = \frac{20}{36}, x_T = 0.9, P(X_{1.2}) = 0.9$ RiskNav Output: *Impact Intensity I = 0.76*
Risk Area 2. This area includes three risk events associated with the system's software development.	Risk Event $X_{2.1}$. Software Size—"Software size exceeds 100K source instructions above initial estimates." Risk Assessment Team Inputs: $x_C = \frac{10}{20}, x_S = \frac{24}{36}, x_T = 0.7, P(X_{2.1}) = 0.8$ RiskNav Output: *Impact Intensity I = 0.62*
Risk Area 2 (Continued)	Risk Event $X_{2.2}$. Spare Memory—"The performance of the system architecture fails to allow the required 50 percent memory reserve." Risk Assessment Team Inputs: $x_C = \frac{1}{20}, x_S = \frac{4}{36}, x_T = 0.2, P(X_{2.2}) = 0.5$ RiskNav Output: *Impact Intensity I = 0.14*
Risk Area 2 (Concluded)	Risk Event $X_{2.3}$. Software Staffing—"Contractor is experiencing a high-turnover rate and fails to meet required staffing levels of 50 full-time development persons." Risk Assessment Team Inputs: $x_C = \frac{2}{20}, x_S = \frac{15}{36}, x_T = 0, P(X_{2.3}) = 0.5$ RiskNav Output: *Impact Intensity I = 0.25*

* These maximums are set by the risk assessment team; for instance, $Max\{x_C\}$ could be the maximum deviation acceptable from the planned program cost; it could also denote the maximum or budgeted program cost.

"induce" risk events $X_{1.1}$ and $X_{1.2}$. Likewise, Figure 6 reveals a weak tendency for risk event $X_{2.1}$, if it occurs, to "exclude" the occurrence of risk event $X_{2.3}$.

When two risk events X_i and X_j, are shown to be moderately to strongly coupled, say

$$L(X_i, X_j) \geq 0.5$$

Table 5. Risk Event Coupling and Coupling Level Determination

Event X_j	Event X_i	$l_{i\|j}$ Eqt. 5	$u_{i\|j}$ Eqt. 6	$P(X_i\|X_j)$ Subjective Assessment	$L(X_i, X_j)$ Rule A or Rule B
2.1	1.1	0.375	0.625	0.60	0.80 Rule A
2.1	1.2	0.875	1.00	0.95	0.50 Rule A
2.1	2.2	0.375	0.625	0.55	0.40 Rule A
2.1	2.3	0.375	0.625	0.40	−0.20 Rule B

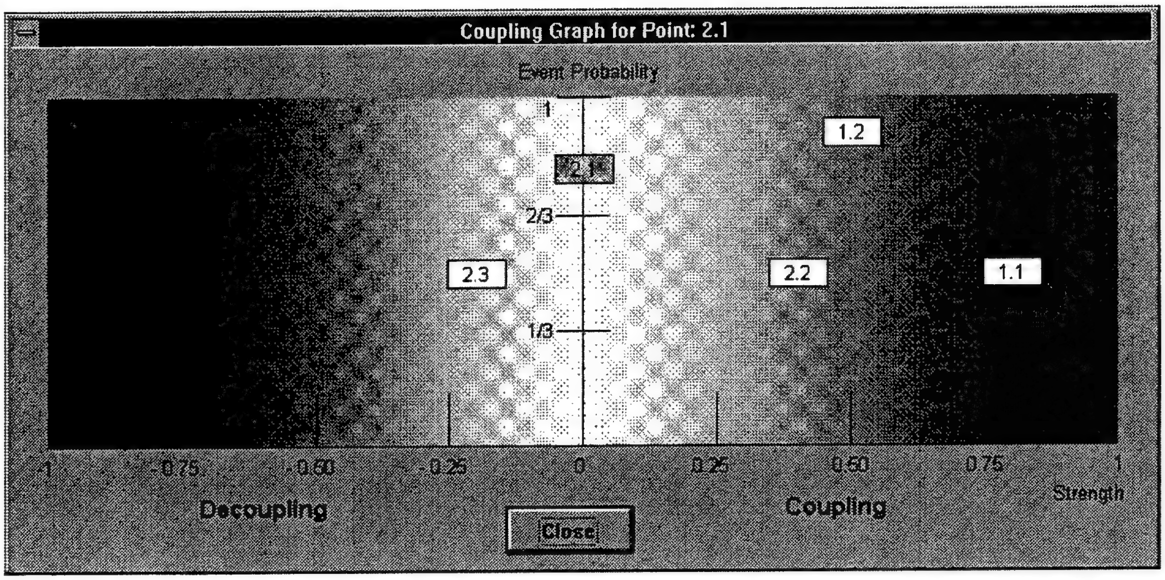

Figure 6. A Risk Event Coupling Display—$L(X_i, X_j)$ from Table 5.

it indicates that they should be considered together in a risk mitigation strategy. Some common root causes between them may be identified so that mitigation strategies can be formulated to reduce their probabilities of occurrence simultaneously.

Note that coupled events do not necessarily imply one causes the other to occur. Occurrences between events that are positively coupled could be due to some cause common to both. Thus, management should not just attack one event to reduce risk and blindly hope for the same result for the coupled event.

When two risk events X_i and X_j are not coupled nor decoupled (or weakly coupled or decoupled), say

$$-0.25 \leq L(X_i, X_j) \leq 0.25$$

their occurrences are independent (or nearly independent) of each other. That is, the probability of the occurrence of one event is not affected by the occurrence or nonoccurrence of the other event. From a risk management perspective, such events should be considered individually in a risk mitigation strategy.

When two risk events X_i and X_j are moderately to strongly decoupled, say

$$L(X_i, X_j) \leq -0.5$$

it indicates a pair that tends toward mutual exclusiveness. The closer $L(X_i, X_j)$ is to negative one the more it becomes that risk events X_i and X_j cannot both occur during the period of interest. That is, both events may not occur but if one of them occurs the other will not occur.

Decoupled risk events also require separate attention by management. Although these events fall into the region where there is a tendency towards mutual exclusion, their individual impact intensities may be significant.

C. Some Remarks

A key concept in RiskNav is its definition of impact intensity. In defining a way to measure this it was decided not to compute the expected impact intensity, given by the formula below

$$E[I(X)] = P(X)I(X)$$

The reason for this is as follows: a risk event with high impact intensity and low probability of occurrence and another event with low impact intensity and high probability of occurrence can produce comparable values for the expected impact intensity. However, they may require different levels of management attention and different risk mitigation strategies. In program risk management it is usually more desirable and/or more cost-effective to focus the risk mitigation effort on reducing events with high impact intensity and high probability of occurrence.

When two risk events are determined to be highly coupled, a natural question to ask next is how they impact the program together if both events occur. In such a circumstance, treat the two events as one and assess the net impact of the combined event on the program in each dimension of cost, schedule, and technical performance. Note that the net impact of the combined event in each dimension may not be just the sum of the individual impacts. They may partially overlap, compensate each other, or have an interactive effect. With the new assessment, the impact intensity can be calculated for the combined event with its probability of occurrence being the joint probability of occurrence of the two individual events. This step can also be extended to a third event and so on.

In the coupling methodology, note that the conditional probability is assessed between two events. From such a pairwise assessment it may turn out that events X_1 and X_2 are positively coupled and so are events X_1 and X_3, but events X_2 and X_3 are determined to be decoupled. Even though this situation is conceivable with certain peculiar event definitions, the contradiction on the surface does indicate possible inconsistency in the conditional probability assessment. The user should take the situation as a warning when it occurs and reexamine the assessments, and/or investigate the events involved, to determine if such a situation is justified.

D. Summary

This paper presents a methodology for prioritizing, displaying, and tracking program risks. It was developed from a program office's need for a straightforward way to isolate key critical risk events from those considered less threatening.

RiskNav embodies a sound risk assessment process, one that can be directly embedded into the systems engineering approach. It's methodology is not mechanically cumbersome; nor is it a highly complex mathematical scheme. RiskNav's prioritization and event coupling displays are very useful in support of risk assessment reviews within a program office or to a program's executive-level decision-makers.

The benefits gained by implementing RiskNav on a program include: the early identification of risk events such that mitigation strategies can be developed in a timely manner, establishing a common set of project-specific cost, schedule, and technical performance scales on which to map risk event impacts, creating a structured environment within the systems engineering process for monitoring and documenting changes in risk events and their prioritizations over time. In the spirit of T. Gilb, RiskNav is one of many possible approaches for "actively attacking risks before they actively attack you (Gilb, 1988)."

Acknowledgments

The authors are grateful to RiskNav's software prototype team; specifically, T. J. Kneeland, D. J. Phair, and G. M. Sawyer of MITRE's Software Prototyping Specialty Group. We also appreciate the independent review of this paper by Z. F. Lansdowne of MITRE's Economic and Decision Analysis Center.

References

AFMCPAM 63-101, Acquisition Risk Management Guide, September, 1993.

Clemen, R. T., *Making Hard Decisions—An Introduction to Decision Analysis*, Duxbury Press, Belmont, California, 1990.

Gilb, T., "*Estimating the Risk*," Chapter 6, *Principles of Software Engineering Management*, Addison-Wesley, Reading, MA., 1988.

Hardy, G. H., Littlewood, J. E., Polya, G., *Inequalities*, Cambridge University Press, 1967.

Keeney, R. L., Raiffa, H., *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, John Wiley & Sons, 1976.

Appendix

A Further Discussion of Event Coupling

In this appendix, a rationale for the rules defined for measuring the coupling level between two events is given along with a discussion of some variations in the rules. As shown in the equations for the coupling level $L(X_i, X_j)$, the measurement rules utilize the joint probability $P(X_i, X_j)$ of risk events X_i and X_j as a base for expressing their tendency to occur together or in each other's absence. As an indicator for such a tendency, however, the joint probability has to be modified as it expresses the likelihood that both events will occur, not exactly the influence that the occurrence of one event has directly or indirectly on the other.

For example, two independent events each with a marginal probability of 0.9 have a joint probability of 0.81. This high probability of both events to occur only reflects the fact that each event has a very high probability to occur, independent of whether the other event is actually to occur or not. Since the two events have no influence on each other's occurrence, their coupling level would be 0. On the other hand, suppose two events each with a marginal probability of 0.1 have a joint probability of 0.1 also. Though low, this joint probability along with the marginal probabilities implies that if any of the two events is to occur, the other is sure to occur. There is a definite "*induction*" of one event's occurrence by the other and thus their coupling level would be $+1$ to indicate such a positive and complete influence. For still another example, assume two events have a marginal probability of 0.5 each and a joint probability of 0. Even though each event has a 50 percent chance to occur, they for sure will not both occur; if one is to occur, the other will not. The influence of one event's occurrence on the other is also complete in this case but it is in a negative sense. The coupling level for these two events would be -1 to indicate such a definite "*exclusion*".

The modification on the joint probability $P(X_i, X_j)$ to produce an appropriate measure for coupling is done in two steps. First, the joint probability is adjusted (subtracted) by a value of $P(X_i)P(X_j)$, which would be the joint probability of the two events should they be independent. Thus, if the adjusted value is 0, the two events must be independent of each other. If it is positive, the two events must be not independent; the higher the adjusted value, the more positively related are (or mutually "*inductive*") the two events' occurrences. On the other hand, if it is negative, the two events must be not independent either; the lower the adjusted value, the more negatively related are (or mutually "*exclusive*") the two events' occurrences.

The value $P(X_i, X_j) - P(X_i)P(X_j)$ is then rescaled (divided) by its largest possible absolute quantity, which is different depending on whether the adjusted value is positive or negative (thus the two rules). The purpose of this step is to make the resulting value for coupling level to vary between 0 and $+1$ for positively related events and between 0 and -1 for negatively related events. When the adjusted value is positive, its largest possible quantity for given $P(X_i)$ and $P(X_j)$ is clearly the value when the joint probability is at its maximum, which can be shown to be the smaller of the two marginal probabilities. Thus, the divider for the adjusted value is $\mathrm{Min}\{P(X_i), P(X_j)\} - P(X_i)P(X_j)$ and that results in Rule A.

When value $P(X_i, X_j) - P(X_i)P(X_j)$ is negative, its largest possible absolute quantity is such when the joint probability is at its minimum. For given $P(X_i)$ and $P(X_j)$, the minimum joint probability can be shown to be $\mathrm{Max}\{0, P(X_i) + P(X_j) - 1\}$. This is to say when $P(X_i) + P(X_j) > 1$, there is a probability at least as large as $P(X_i) + P(X_j) - 1$ that the two events will both occur, i.e., they cannot be totally exclusive of each other. To reserve the coupling level of

-1 for the case when two events are totally exclusive of each other (i.e., $P(X_i, X_j) = 0$), we purposely use 0 in place of $\text{Max}\{0, P(X_i) + P(X_j) - 1\}$ to obtain the largest possible absolute quantity $P(X_i)P(X_j)$. Dividing the negative adjusted value by this quantity gives Rule B.

Note from the rules $L(X_i, X_j) = L(X_j, X_i)$, i.e., the coupling measure we define here is the same for either ordering of the events, much as the case for the joint probability $P(X_i, X_j)$. This symmetrical property of the coupling level seems consistent with events independent or totally exclusive of each other as either relationship can be shown to be mutual between the two events involved. However, when two events are such that one definitely *"induces"* the other, the reverse may or may not be true. This may seem inconsistent with the symmetrical property of the coupling level (+1 in this case). One way to make the coupling level $L(X_i, X_j)$ to respond to this non-mutual relationship is to replace the term $\text{Min}\{P(X_i), P(X_j)\}$ in Rule A by $P(X_j)$. With this change, $L(X_i, X_j) = +1$ when $P(X_i|X_j) = 1$, but $L(X_j, X_i) \le +1$ with the equality held true only if $P(X_j|X_i) = 1$ (i.e., the coupling level is symmetrical if the definite *"induction"* relationship is mutual). But, this change would make the coupling level non-symmetrical in general and would require more calculation for each pair of positively related events. We decided in this paper not to adopt the change but to select symmetry for simplicity.

Although this paper has presented one approach to measuring risk event coupling, other ways are possible. For example, the correlation between two risk events can be used to measure coupling, as one referee suggested. Specifically, define a Bernoulli random variable for the event: $Y_i = 1$ if X_i occurs and $Y_i = 0$ otherwise. We then have

$$E(Y_i) = P(X_i)$$

$$\text{Var}(Y_i) = P(X_i)[1 - P(X_i)]$$

$$\text{Cov}(Y_i, Y_j) = P(X_i, X_j) - P(X_i)P(X_j)$$

From the above relationships, the correlation between Y_i and Y_j becomes:

$$\text{Corr}(Y_i, Y_j) = \left[\frac{P(X_j)}{1 - P(X_j)}\right]^{1/2} \frac{[P(X_i|X_j) - P(X_i)]}{[P(X_i)\{1 - P(X_i)\}]^{1/2}}$$

If the risk events are independent, then this measure is zero because the difference in the numerator of the second term is zero. If X_j induces X_i, so that $P(X_i|X_j) = 1$, and $P(X_j) \le P(X_i)$, then

$$\text{Corr}(Y_i, Y_j) = \left[\frac{P(X_j)}{P(X_i)}\right]^{1/2}\left[\frac{1 - P(X_i)}{1 - P(X_j)}\right]^{1/2} \le 1$$

In fact, the bound of 1 is achieved only if $P(X_j) = P(X_i)$. If X_j excludes X_i, so that $P(X_i|X_j) = 0$ and $P(X_i) + P(X_j) \le 1$, then

$$\text{Corr}(Y_i, Y_j) = -\left[\frac{P(X_j)}{1 - P(X_j)}\right]^{1/2}\left[\frac{P(X_i)}{1 - P(X_i)}\right]^{1/2} \ge -1$$

This bound is achieved only if $P(X_i) + P(X_j) = 1$.

The measure proposed in this paper, $L(X_i, X_j)$, has the advantage that the bounds of ± 1 are always achievable, regardless of the marginal probabilities. However, the correlation approach described above has an easy interpretation.

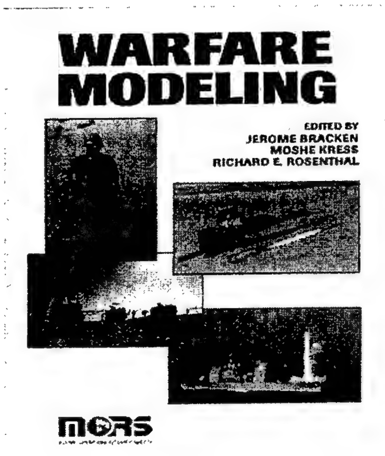

ORDER FORM

Warfare Modeling

Published by the
MILITARY OPERATIONS RESEARCH SOCIETY
In cooperation with John Wiley & Sons
600 pages, $35.00.

The Military Operations Research Society (MORS) serves the students, developers, practitioners, and clients of military operations research through the ongoing production of high-quality, low-cost meetings, books, and periodicals.

Warfare Modeling was selected for publication by MORS because it contains 25 state-of-the-art chapters, addressing a variety of operations research models that are of central importance in military planning, analysis, and operations. Research on this subject matter seldom appears in the open literature. Most of it appears in reports and documents of government agencies and advisory corporations.

This book is an invaluable reference for military OR professionals as well as for a more general audience of researchers and practitioners. The scope of models included (analytic and simulative, stochastic and deterministic, simple and complex, applied and theoretical, domestic and international) is extensive.

The Book Is Comprised Of Seven Sections:

1. Theater-Level Modeling
2. Mathematical Models of Combat
3. Historical Analysis
4. Weapon System Analysis
5. Command, Control, Communications, and Intelligence
6. Cost Effective Analysis
7. Modeling in Support of Operations Desert Shield and Desert Storm

Warfare Modeling was edited by Jerome Bracken of Yale University, Moshe Kress of Israel's Center for Military Analyses, and Richard E. Rosenthal of the US Naval Postgraduate School. The Forward is by Wayne P. Hughes, Jr., a retired Navy captain and MORS Fellow.

Please Send Me _____ copies of *Warfare Modeling* at $35 per copy

plus $5 for US freight, $8 for Canada, and $10 for all other international shipments per book.

Name ____

Address ______________________________________

____

____

City ____________________ State __________ Zip Code ________

VA Residents add $1.58 sales tax per copy

☐ CASH ☐ CHECK ☐ Visa ☐ MC ☐ AMEX Credit card # ____________________ Expiration date: __________ Billing zip code: __________ Name on Credit Card (print): ____________________ Phone Number: ____________________ Signature: ____________________	TOTAL $

101 South Whiting Street • Suite 202 • Alexandria, VA 22304 • (703) 751-7290 • Fax (703) 751-8171
E-mail morsoffice@aol.com

The Probabilistic Multiple-Travelling-Salesmen Facility-Location Problem: Space-Filling Curves and Asymptotic Euclidean Analyses

Yupo Chan[1]

Air Force Institute of Technology

David L. Merrill

Air Mobility Command

ABSTRACT

The Probabilistic Multiple-Travelling-Salesmen Facility-Location Problem locates a depot and constructs the *a priori* tours for visiting all demand instances at n nodes. It is an NP-hard problem that defies straightforward solutions. We offer an $O(n^2)$ solution algorithm based on Space-Filling-Curve (SFC) heuristics. A companion $O(n)$ algorithm suggests that the 1-median and 1-center solutions can be good approximations for such a facility-location problem. A flight-inspection problem in the United States Air Force serves as an ideal case-study for the methodology proposed here. Even though a solid recommendation cannot be offered to the inspection problem due to insufficient data, the case study allows us to illustrate the methodology with clarity. For example, a worst case error bound for the heuristics ensures the solution is within $O(\log n)$ of the optimum. By comparing with linear-programming relaxation solutions to 52 instances of the problem, the heuristic solution is shown to be within 7.6 percent of the optimal on the average. Short cuts for locating depot is suggested using the 1-center and 1-median approaches, where route lengths are approximated by Euclidean distances. The 1-center and 1-median approximations are validated based on asymptotic properties and computational experiences with the top 50 Air-Force flight-check bases. Finally, we illustrate how the methodology can offer operators in the field a simple but effective means to respond to changing demands on a day-to-day bases with trivial computational requirements.

I. INTRODUCTION

In the last decade, there has been an explosion of activities in location and routing analyses, particularly combined location-routing models (Chan forthcoming). Even more exciting is the probabilistic formulation of such class of problems (Laporte et al. 1987, Dror et al. 1989, Bertsimas et al. 1990, Bertsimas and Van Ryzin 1991, Bertsimas 1992, Bramel 1992, Bramel et al. 1992, Jamil et al. 1994, Laporte et al. 1994, Bertsimas and Simchi-Levi 1996). These activities are motivated by both practical applications and the intellectual content of such problems. The requirement for just-in-time deliveries, for example, is urgent in the industrial sector as much as emergency responses in the defense community. The exciting computational advances go hand-in-hand with research in complexity theory and "duality" in integer programming.

This paper documents research in a Probabilistic Multiple-Travelling-Salesmen Facility-Location Problem (PMTSFLP). The problem is to construct k *a priori* tours and to pre-select a location for the depot that will minimize the expected cost of all instances of the demands. A natural problem arises in the United States Air Force regarding its flight-inspection missions to all Air Force Bases (and selected civilian airports.) Here, demands vary from week-to-week among the bases. The acquisition of a brand new inspection-aircraft fleet also motivated the re-examination of basing decision. Finally, the lack of analysis capability at the Squadron level dictated some *a priori* scheme that can easily be modified for special instances in the field without extensive computational requirements.

While the literature stopped short at Probabilistic Travelling-*Salesman* Facility-Location Problems, this research extends the findings to multiple *salesmen*. The Space-Filling-Curve (SFC) heuristic is also extended to provide the implementation tool. Because of the computational efficiency of SFC, the combined location-routing problem can be solved by *re-optimization,* wherein the relative merits of each basing decision can be easily updated from a previous solution and then compared with one another. *For a changed problem instance—whether due to different demands or fleet size—re-optimization can be readily performed.* The "best" can then be selected, commensurate with its efficient location-routing strategies. Finally, the SFC heuristic allows for quick response to a particular instance of the problem with practically trivial computational requirements, thus allowing operators in the field to respond quickly to changing conditions on a day-to-day basis.

In this paper, we first formulate the location-routing model, particularly its probabilistic version. This formulation is, to the best of our knowledge, the first of its kind and not found in the literature. This serves as the lower bound for validation

APPLICATION AREAS:
Mobility, Logistics

OR METHODOLOGIES:
Facility Location,
Probabilistic Travelling Salesmen,
Space-filling-curve,
1-Median, 1-Center

[1] Correspondent is professor, Dept. of Operational Sciences, 2950 P Street, Air Force Institute of Technology, Wright-Patterson AFB, Ohio 45433-7765.

This paper won the 1991 Koopman Prize of the Military Applications Society of the Institute of Operations Research and Management Science

purposes. The model is used to measure the performance of a fast SFC heuristic: the thrust of this research. The results of a carefully-designed case-study are validated not only by the exact mathematical-programming model, but also the approximate *median* and *center* solutions. These validations, including analytic error bounds, all point toward the viability of solving the PMTSFLP using the proposed SFC approach.

II. LOCATION-ROUTING MODELS

Selecting a location for basing a fleet of inspection-aircraft can be modelled by several means. If aircraft routing is not of concern, a median or center formulation may be sufficient. On the other hand, a combined routing-and-location model becomes mandatory when a base is to be located specifically such that the amount of actual flying is kept at a minimum. By way of a definition, the 1-median problem seeks to locate a facility so that it is proximal to all demands on the average. Let us use p_i to denote the probability of demand at node i, and d_{ij} as the Euclidean distance from demand i to facility j, y_j as the binary variable to locate facility at site j, and x_{ij} as the binary variable indicating that the facility at site j will serve demands at site i. The 1-median problem where I is the set of n nodes in a network is simply a special case of the p-median problem when $p = 1$. This special case can be solved very simply by evaluating the objective-function value for locating the single facility at each of the n nodes. The location that yields the smallest value is then selected. The 1-center is defined as the location where the most distant demand node from it is as close as possible and can be computed equally with ease (Daskin 1995).

A. Simplified Analytic Model. The classic Travelling-Salesman-Problem (TSP) seeks a minimum tour among all n nodes (Hoffman and Wolfe 1985). The more general binary-programming-formulation for Multiple Salesmen, or multiple-TSP, can be stated as

$$\min_x \sum_{i\in I} \sum_{j\in I} d_{ij}x_{ij} \tag{1}$$

$$s.t. \sum_{i\in I} x_{ij} = \begin{cases} k & \text{if } j = 1 (\text{or the depot}) \\ 1 & \text{if } j = 2, 3, \dots, n. \end{cases} \tag{2}$$

$$\sum_{j\in I} x_{ij} = \begin{cases} k & \text{if } i = 1 (\text{or the depot}) \\ 1 & \text{if } i = 2, 3, \dots, n. \end{cases} \tag{3}$$

where $x_{ii} = 0$, $x_{ij} \in \{0,1\}$, and k is a given integer ($1 \le k \le K$) specifying the number of aircraft deployed in the fleet. Fleet sizes ranged between 2 and 4 in our case study (as will be shown in Section V). In addition, subtour-breaking constraints should be implemented.

A PMTSFLP is characterized not only by uncertain demand, the number of "salespersons" (airplanes) sent "on the road" k is also a variable, only limited by the fleet availability K at the depot. The actual number of airplanes used to solve the flight-inspection problem is determined by the squadron commander based on timeliness-of-response and a "reasonable" time for the flight crew to be "on the road." For a depot based at node 1, we have

$$K \ge \sum_{j\in I} x_{1j} \ge 1 \tag{4}$$

$$K \ge \sum_{i\in I} x_{i1} \ge 1 \tag{5}$$

For small networks of seven nodes or less, the set of linear (rather than binary) programming constraints, (4) through (7), works extremely well in replacing Equations (2) through (3) above, including the subtour-breaking constraint.

$$x_{ij} + x_{ji} \le 1 \quad \text{for all } i, j \text{ not equal to } 1 \tag{6}$$

$$x_{i1} + x_{1i} \le 2 \quad \text{for all } i. \tag{7}$$

The above states that a trip must enter and leave a demand node via different ways except to and from the depot. In other words, these two constraints effectively keep aircraft from "doubling back" except to and from the depot, which means the minimum subtour would have to involve three bases. Since our case study initially involves low bases-visited-to-aircraft ratio, this effectively precludes subtours. Finally,

$$\sum_{i\in I} x_{ij} = 1 \quad \text{for all } j \text{ except the depot } (i \ne j) \tag{8}$$

$$\sum_{i \in I} x_{ji} = 1 \quad \text{for all } j \text{ except the depot } (i \neq j) \tag{9}$$

This constraint forces exactly one entry and one exit to and from every node other than the depot. Again, the above simplified multiple-TSP linear program only works up to seven nodes (including the depot) in our experimentation.

B. Exact Analytic Model. To ensure that these simplified subtour-breaking constraints work, "conventional" subtour-breaking constraints are used to validate the resulting tours:

$$\delta_i - \delta_j + |I| x_{ij} \leq |I| - 1 \qquad 2 \leq i \neq j \leq |I| \tag{10}$$

where δ's are real numbers that count the number of "hops." As a final check, crew-duty-day restrictions are enforced. This is to reinforce our concept that the tours are designed, both manually by the commander in the field as well as the analytical models here, to fit a mission (i.e. each salesman-tour) within the allowable "time on the road":

$$\sigma_i - \sigma_j + U x_{ij} \leq U - d_{ij} \quad \forall 2 \leq i \neq j \leq |I| \tag{11}$$

where σ_i's are the real variables representing an "odometer reading" at node i, recording the amount of travel up to node i on the tour. Here, U is the crew-duty-day limitation in hours. The above subtour-breaking equation (11) ignores the leg to and from the depot. It is clear that the σ associated with the first demand point on a tour should equal the distance to the home depot. Additionally, the distance from the last node prior to returning to the depot cannot exceed the remaining "range" U of the vehicle. The following additional equations (Chan and Baker 1996) sufficiently restrict range to complete equation (11).

$$\begin{array}{ll} d_{ij} x_{ij} \leq \sigma_j & i = 1, 2 \leq j \leq |I| \\ d_{ji} x_{ji} + \sigma_j \leq U & i = 1, 2 \leq j \leq |I|. \end{array} \tag{12}$$

In all cases, the simplified constraints and the crew-duty-day requirements are validated (Reynolds et al. 1990). While not a proof for the formulation, Equations (4)–(9) appear to work well to break subtours for small networks.

The combined location-routing model, the main subject of this paper, is the combination of the 1-median and multiple-TSP where d_{ij} in Equation (1) is replaced by D_{ij} to mean distance measured along a tour that connects depot j to demand i, rather than a Euclidean distance between i and j necessarily. Here the distance measure $D_{ij} = \sigma_i - \sigma_j$ is the difference in "odometer readings," where the "odometer" records travel mileage on the tour. For the *forcing/linking* constraint between the location and routing problems, one can write (Loftus and Chan 1992).

$$\sum_{j \in M_i} x_{ij} \leq K y_i \quad \forall i \in J \tag{13}$$

$$1 \leq \sum_{i \in N_j} x_{ij} \leq K y_j \quad \forall j \in J \tag{14}$$

where M_i is the set of nodes emanating from node i and N_j the nodes incident upon j. These constraints are somewhat tricky. For example, when $K = 2$, the terms Ky_i and Ky_j need to be changed to $1\text{-}y_i$ and $1\text{-}y_j$ respectively when demand exists at a depot. Also we write for the set of candidate-depot sites J.

$$x_{ij} + x_{ji} \leq 1 + \sum_{k \in J} y_k \quad \forall i \in J, \quad \forall j \in J \tag{15}$$

$$x_{il} + x_{li} \leq 1 + y_1 \quad \forall\, i \in I - J, \quad \forall l \in J \tag{16}$$

which establishes the precedence relationship between basing decisions y and routing decisions x. Again, these constraints are tricky. For K=2 and when demands exist at the depots, for example, Equation (15) is no longer required (Shirley and Chan 1993). In lieu of an explicit linking/forcing constraint, however, Merrill (1989) enumerated all possible depots that can base a vehicle fleet and pick the one with the lowest cost. His approach essentially uses a sequential approach to solve location-routing problems. It is justifiable since all depots are existing depots and they are few in number; little or no additional costs will be necessary to base a vehicle fleet there. Thus a location-routing model degenerates into several routing models based at candidate depot sites. Another view of this is to consider the 1-median problem, where distances are in the subgraph defined by the k tours of the multiple-TSP.

C. Summary. This concludes our analytic formulation of PMTSFLP. Because of the probabilistic nature of the demand, *a priori* optimization is employed here, whereby a depot is to be located in anticipation of future demands. Another characteristic of the current formulation is that a *variable* number of vehicles can be dispatched, limited only by the fleet size available at the depot *in response to a demand instance.* The analytic formulation advanced in this section includes the simplified formulation for small networks (up to seven nodes), consisting of Equations (1) and (4) through (9). The simplified model is a linear-programming model that has been shown to work by an exact, mixed-integer-programming formulation consisting of Equations (1)–(3) and (10)–(16). In the exact model, care is taken to distinguish between the case where demand exists at a depot and where it does not. The simplified model, owing to the small-size network and a preprocessing of missions (see Table I), does not need to employ a "range" constraint to limit the time a vehicle spends "on the road." The exact formulation, on the other hand, does—since the exact formulation is designed to handle large size networks covering a large geographic area in general. Most importantly, these analytic models—the LP formulation in particular—provides the lower bound for validating the a-priori-optimization heuristic, corresponding to a *fixed* number of salespersons (aircraft) as shown in Table I. To the best of our knowledge, both analytic formulations are the first of their kind and have not been reported elsewhere in the literature. While the LP formulation is fresh and practical for our problem, the combined location-routing models adds to the existing knowledge base, particularly the introduction of Equations (12) through (16).

III. PROBABILISTIC MULTIPLE-TRAVELLING-SALESMAN FACILITY-LOCATION PROBLEM

The Probabilistic Travelling-*Salesman* Facility-Location Problem (PTSFLP) can now be formally defined as a location-routing problem—as stated in section II—in which the demands are random and the number of salespersons is a variable. Thus demand comes from node i with probability p_i. The demand list constitutes the basis for planning an optimal tour and locating the depot, so that the expected tour-length is at its minimum. Thus if demand at node i exists then the node must be visited; otherwise, a visit is not required. The magnitude of the demand is not relevant in our formulation; therefore, demand is modelled as a random variable with realizations zero and one. Obviously, one way to solve the PTSFLP is to solve for all the instances of the constituent Travelling-Salesman Facility-Location Problem (TSFLP), even though this can be an extremely

Table I. Selected Database of Flight Inspections from Scott

Mission	No. of aircraft	Bases visited					
1	2	Moody	Dobbins	MacDill	Tyndall	Hurlburt	
2	3	Patrick	MacDill	Kingsville	Ellsworth	Range 1	
3	2	Williams	March	Moody			
4	3	Barksdale	Oklahoma City	Dallas	Ellington	Kingsville	
5	2	Williams	McClellan	McChord	Andrews	England	
6	2	Libby	Davis-Monthan	Reese	Columbus		
7	3	Kessler	Patrick	Volk	Richards-Gebaur		
8	3	Offutt	Little Rock	Robins	Myrtle Beach	Shaw	Moody
9	4	Homestead	Palmerola	Andrews	Kirtland	Offutt	Range 1
10	3	Richards-Gebaur	Oklahoma City	Fort Smith	Barksdale	Redstone	
11	2	Cannon	Amarillo	Oklahoma City	Kirtland	Nellis	Range 2
12	4	Little-Rock	Wright-Patterson	Carswell	Nellis	Peterson	Langley
13	3	Pope	Richards-Gebaur	Fairchild	Carswell	Ellsworth	Maxwell

laborious process. But there are some recent findings on the TSFLP that can be exploited to bypass this imposing (if not impossible) task. We will generalize these observations to the Multiple-Salesmen case ($k > 1$), as formulated in PMTSFLP.

Distinction is made here between *a priori* tours, where it is possible to minimize the expected tour-length by taking the expected-value of the random-demand vector, and *posterior* tours, where one selects a minimum-length tour for each demand-realization by complete enumeration. When the number of aircraft k is fixed in the optimization, lengthy tours may result, which can be longer than the squadron commander likes to have the crew away from their family. *Post-optimization* caters for these situations when *optimization* violates the "time on the road" constraint (11). More aircraft are subsequently used to shorten crew-duty days. As k is increased to satisfy the "time on the road" constraint for each salesperson (aircraft crew), we pay for this through an increased total tour-length among all salespersons (i.e. inducing feasibility for a "super-optimum"). The total tour-length is monotonically non-decreasing as we move toward a 1-median depot solution (or when $k \to n$). As we approach the 1-median solution, the total tour-length is at its longest while individual salesman-tour is at its shortest. Obviously, *post-optimization* is of secondary importance here, since it is relatively straightforward. It is used in this paper to conduct control experiments, in which historical flight records characterized by reasonable crew-duty days are revisited to assess how historical performance measures against optimization.

A. Theory. Hakimi (1964) suggested that the medians of a network are found among the nodes. Berman and Simchi-Levi (1988) extended Hakimi's original finding to the TSFLP, suggesting that the optimal location has to occur at a node also. The PTSFLP then reduces to finding node j for the expected tour-length of TSFLP that minimizes $v(j) = \Sigma_{S \subset I} P(S)\, L(S \cup j)$. Here $P(S)$ is the probability that the instance occurs with nodal demands realized at a subset of nodes S in I. $L(S \cup j)$ is the length of the optimal Travelling-Salesman-Tour (TST), $T(S \cup j)$, for the instance S with the base located at j. And for independent p_i's, $P(S) = \Pi_{i \in S}\, p_i\, \Pi_{i \in I-S}\, (1 - p_i)$. Bertsimas (1989) showed that if $p_j = 1$, or demand exists at a facility with certainty, the optimal location for the facility (depot) is in fact at j if the distance matrix satisfies the triangle inequality $d_{ij} \leq d_{ik} + d_{kj}$. Furthermore he extended a solution heuristic by Berman and Simchi-Levi (1988a), wherein a relative worst case error of $(1 - p_j)/2$ is achieved, when j is the optimal location. When $p_j = 1$ for all j, the problem reduces to a deterministic TSFLP, and the solution is expected to be exact.

As suggested previously, it is cumbersome (if not impossible) to compute the optimal tour in every instance. In the Probabilistic Travelling-Salesman Problem (PTSP) literature such as that by Berman and Simchi-Levi (1988a) and Laporte et al. (1994), the concept of an *a priori* tour based at j is introduced, $T(j)$, visiting all *potential* demands. For a particular instance, one skips nodes without a demand. The problem then boils down to finding the location j and the "master" tour T that minimizes $g(j, T) = \Sigma_{S \subseteq I} P(S)\, L_T\, (S \cup j)$ where $L_T(X)$ is the length of the master tour restricted to sites in X. This brings us face-to-face with the PTSFLP, where it has been shown that the optimal location occurs at a node also (Berman and Simchi-Levi 1988a).

B. Algorithm. The PTSFLP can be reduced to the solution of n PTSPs. Given the vector of probabilities $(p_1, p_2, \ldots, p_n)$ and the optimal location at node j, the corresponding optimal tour T_j is the Probabilistic Travelling-Salesman Tour (PTST) with the vector of probabilities $P_j = (p_1, \ldots, p_{j-1}, 1, p_{j+1}, \ldots, p_n)$, where the depot carries with it a unitary probability. The problem then boils down to finding the n optimal PTSPs corresponding to the vectors of probabilities P_j $(j = 1, 2, \ldots, n)$, and then select the one with the minimum expectation. The resulting j^* is the optimal location. It can be shown that this result can be generalized to the multiple-salesmen case as in PMTSFLP. Independent of one another, Bertsimas (1988) and Merrill (1989) arrived at a similar SFC heuristic for solving the PTSFLP and the PMTSFLP respectively. Here, we state an extension of Merrill's solution, which is a more general algorithm for the *multiple*-salesmen application to be discussed in this paper:

Step 1. Given the coordinates of the loca-

tions of all the demands in Set I, use SFC to find *a priori* master-tours T_k or the tours of k salespersons. (See detailed algorithm for SFC in the following section.)

Step 2. Compute for every j $(j = 1, 2, \ldots n)$, $G_n(j, T_k)$, where G_n is the generalization of the tour-length function g with k vectors of probabilities $P_j(k)$ corresponding to k salespersons covering n demand nodes without replication. Specifically $P_j(k) = (p_1, \ldots, p_{m_k}; p_j = 1)$, where $m_k \leq n$, and with the set of nodes I partitioned into $I = I_1 \cup I_2 \cup \ldots \cup I_k$, where $I_i \cap I_j = \varnothing$ for $i \neq j$. The union notation simply denotes the partition of n nodes on the space-filling curve into $m_1, m_2, \ldots, m_k$ number of nodes based on clustering, with $m_1 + m_2 + \ldots + m_k = n$. An extended SFC heuristic, called multiple-SFC is used to operationalize this, as will be explained in sequel.

Step 3. Select the point j^* $(j \in I)$ that minimizes $G_n(j, T_k)$. Location j^* and the tours T_k^* constitute the solution to PMTSFLP, call this $G_n^*(j^*, T_k^*)$. If *optimization* is sought where k^* is given, stop; otherwise proceed.

Step 4. For *post-optimization* with all potential demand instances $S \subseteq I$ and where k is a variable, perform steps 1 through 3 for all k where $k^* \leq k \leq K$. The minimum master tour-length $G_n^{**}(j^{**}, T_k^{**}) = \min_{k^* \leq k \leq K} G_n^*(j^*, T_k^*)$ determines the optimal facility-location j^{**}, the optimal tours $T_{k^{**}}^{**}$, and the number of vehicles k^{**} to be employed to serve all potential demands. □

Aside from the optimization vice post-optimization distinction, *Step 4* is necessary in view of the "time on the road" constraint (11), which could make T_k^* infeasible. Inasmuch as SFC does not have a range constraint explicitly built in, *Step 4* will enforce the "range constraint," assuming K is large enough to make it possible.

IV. SPACE-FILLING-CURVE HEURISTICS

Space-filling-curves (SFC) constitute an extremely fast heuristic to solve the planar TSP. It is an $O(n \log n)$ method based on sorting and works as follows (Bartholdi and Platzman 1988):

Step 1. Given the n coordinates (X_i, Y_i) of the nodes in the plane, compute the number $f(X_i, Y_i)$ for each node. The function f is called the Sierpinski curve.

Step 2. Sort the numbers f in ascending order and visit the nodes in the same order, producing a tour T. (It is convenient to have f ranged between 0 and 1.) □

A. Properties. Consider an instance S of the problem. Suppose the SFC heuristic produces a tour $T(S)$ if we run the heuristic on the instance S representing a realization of demands at a subset of nodes I. Let $T(0)$ be the tour produced by the heuristic on the original instance. Then $T(S) = T(0)$, or the order of visit among the two tours are the same according to step 2 of the SFC heuristic above. The reason is that the sorting preserves the order, which is exactly the property of the probabilistic travelling-salesman-problem (PTSP) as well (Bertsimas 1989). This is akin to the observation by Hargrave and Nemhauser (1962) that the order of visitation around the convex hull circumscribing a network does not change in reaching the optimal solution. If one can map points in Euclidean space onto a unit square and overlay a circuit generated by a SFC which visits all nodes within the square (as shown in Figure 1), then the order which the nodes are visited as one travels along the curve will approximate the order the points would be visited by a travelling salesman on an optimal tour. Closeness of orders in this case refers to travel length on the tour.

Bartholdi and Platzman (1988) went on to show that if two nodes are close on the SFC, then they are close on the plane as illustrated by points i and j in Figure 1. Conversely, if two nodes are close in the plane, then they are *likely* to be close on the curve. In the example shown in Figure 1, it is true for points $i - j$, but not for $i' - j$. In other words, i is close to j both on the plane and on the SFC. Even though i' is close to j physically on the plane, however, they are not close on the SFC (at least not as close as the i–j pair anyhow). A set of n nodes projected into a square of area A and mapped onto a SFC will always result in a tour length no greater than $2\sqrt{nA}$, which is also the limit for the TSP tour. For a unitary square, this reduces to $2\sqrt{n}$. The ratio of heuristic tour-length to theoretical is

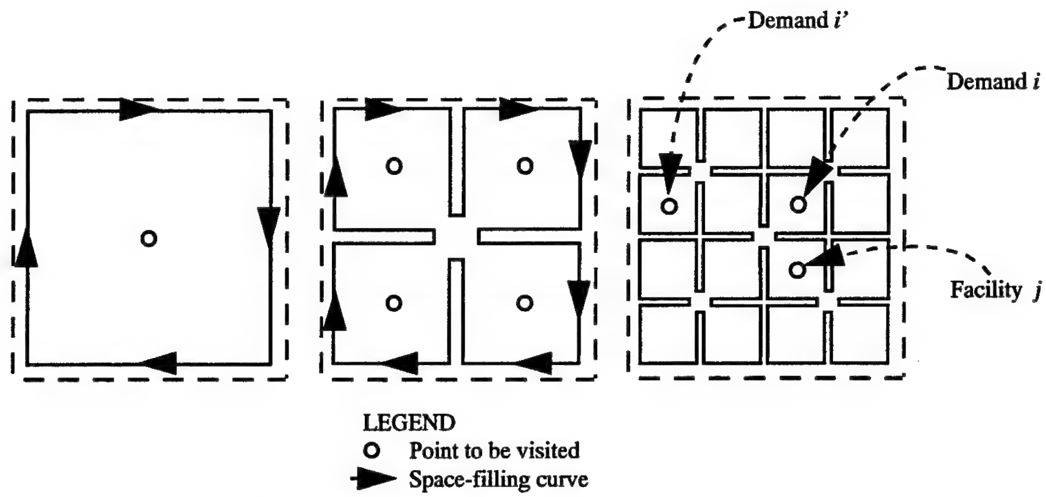

Figure 1 - "Space-filling curve 'visits' (encircles) all points in a square"

$O(\log n)$. Asymptotically, the heuristic tour will be at most 25 percent greater than an optimal tour for very large number of random demand-points ($n \to \infty$). The algorithm is extremely fast, consisting of essentially sorting. Its computational complexity is $O(n \log n)$ in the worst case, and $O(n)$ in the expected case. In the terms of Bartholdi and Platzman (1988), the heuristic is *abstemious* in its data requirements. Only the $O(n)$ coordinates of the points to be visited are necessary, while the $O(n^2)$ distances between points are ignored. In other words, information on the metric and the distribution from which the points are drawn are not required. The algorithm is *agile* in that demand points may be inserted into or deleted from the tour and the solution can be *updated* within $O(\log n)$ steps. Finally, the computer coding is trivial.

It is possible to solve certain multiple-travelling-salesmen problems (MTSP) by partitioning routing tasks (Bartholdi et al. 1983). Demand nodes are assigned to one of the salesmen, who in turn performs TSPs on his assigned nodes. An example was mentioned (Bartholdi and Platzman 1988) in which demand nodes are uniformly distributed among a square. By partitioning the TSP heuristic into segments containing equal number of nodes, a solution is obtained corresponding to segments of roughly equal lengths. Extending the SFC heuristic to handle the multiple-salesmen problem through the preservation of nearness of data points for each salesman, we obtained a solution to the MTSP which will be illustrated in our case study.

The multiple-SFC heuristic outlined in Section III.B to solve the PMTSFLP can be executed in $O(n^2)$ time. Step 1 of the algorithm can be executed in $O(n \log n)$ time since it is nothing more than a straight application of the SFC heuristic with partitioning introduced. Now notice the difference in calculating $G_n(j, T_k)$ and $G_n(j+1, T_k)$ in step 2 is only in the probability vectors, which differ merely in the jth and $(j+1)$th positions. For this reason, we can update from $G_n(j, T_k)$ in $O(n)$ time during reoptimization. After the computation for $G(1, T_k)$ in $O(n^2)$ time, the rest can be completed in $O(n)$. The net computational complexity is $O(n^2)$. Step 3 clearly takes $O(n)$ time, and step 4 $O(kn^2)$, with $k \ll n$. The final result is therefore $O(kn^2)$ for the entire algorithm.

The PMTSFLP heuristic yields a tour length at a factor of $O(\log n)$ from the optimal PMTSFLP-tour-length [Bartholdi and Platzman (1988, 1982)]. It is asymptotically optimal and equivalent to the Multiple-Travelling-Salesmen Facility-Location Problem (MTSFLP). In other words, Bartholdi and Platzman (1988) proved that $g(i, T)/v(i) = O(\log n)$ where g and v are the SFC-heuristic and analytical solution to the PTSP through re-optimization. The MSFC clustering partitioning for PMTSP is within a constant factor r of the best, since proximity on the Sierpinski curve means proximity on the plane. To the extent that proximity in the plane means only a likely proximity on the curve, it is expected that $r \geq 1$. The conjecture is the $r = 1 + e$ where e is a small positive number. Hence $G_n^{**}(j^{**}, T_k^{**})/w(j^{**}) \leq r\, O(\log n)$ in the PMSFLP algorithm in Section III.B, where $w(j^{**})$ is the optimal tour-length from the analytic model and r is likely to be a problem-specific constant. Combining the results from the PMSFLP algorithm above, and for n demand points that are randomly distributed.

$$\begin{aligned} G_n(j, T_k)/w(j^{**}) &\leq G_n^{*}(j^*, T_k^{*})/w(j^{**}) \\ &\leq G_n^{**}(j^{**}, T_k^{**})/w(j^{**}) \\ &= O(\log n) \qquad (17) \end{aligned}$$

where r is problem dependent and $r \leq 1.25$. In the asymptotic case when n is sufficiently large, $\lim_{n\to\infty} \sup G_n(j^{**}, T_k^{**})/w(j^{**}) = 1.25$.

B. Example. To illustrate the simplicity of the multiple-SFC heuristic, we have included here in Figure 2 a graphical illustration of Mission 13, which involves three aircraft and seven

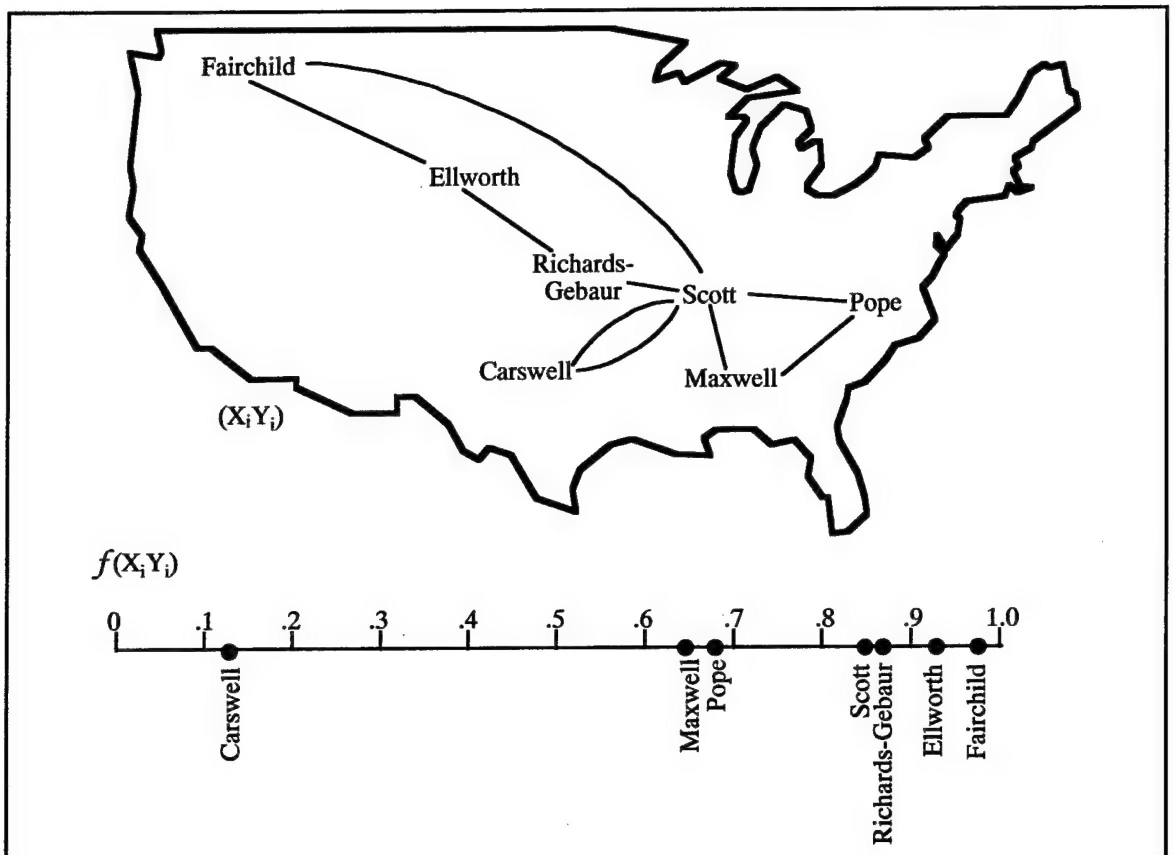

Figure 2 - "The multiple-space-filling-curve heuristic"

bases (including the home depot of Scott Air Force Base). We show the transformation of the (X, Y) coordinates of the seven bases into the Sierpinski-curve function $f(X, Y)$. These f-values allow a ranking of the bases between the scale of 0 to 1. Clusters of these bases suggest multiple-TSP tours. Using three aircraft, the logical tours constructed out of the clusters are:

Aircraft 1: Scott - Carswell - Scott
Aircraft 2: Scott - Maxwell - Pope - Scott
Aircraft 3: Scott - Richards-Gebaur - Ellsworth - Fairchild - Scott

In forming these clusters, remember the Sierpinski curve is actually continuous, with the end points 0 and 1.0 being the same.

We like to emphasize that Figure 2 is simply an extract from a much larger Sierpinski curve, consisting of 98 Air Force bases—a problem that will surpass present computational facilities should an analytical solution be attempted. Showing the other 91 bases on the curve is possible but impracticable, and it will confuse the point we are trying to illustrate here. But the principle is the same. In general, the k intervals (tours T_k) in the Sierpinski curve can be chosen in a number of ways (Bowerman and Calamai 1994). The first is the naive approach of breaking the curve up into intervals of n/k points. The second is to break the curve up at the largest gaps. The third is to break the curve at k sites where the demands are all found within a small circle of each other. We suggest breaking the curve up at the k largest gaps, for preservation of proximity and minimization of individual-salesman tour-lengths. These clusters will form the demands which tours $T_k(j)$ will be constructed from a candidate home-base j.

C. Optimality vs. Post-Optimality Partitioning. Notice that when k is determined this way, or by the LP-relaxation or mixed-integer-programming formulation of Equations (1) through (16) it is a random variable. Thus in the SFC heuristic, if all tours are selected *a priori*, and the partitioning into individual tours for each salesman is performed *after* observing the potential demands, then the number of salespersons is a random variable. This is captured in *Steps 3* and *4* of the PMTSFLP algorithm in Section III.B. On the other hand, if all decisions are made based on when k is specified with the data in Table I and in accordance with the LP-relaxation formulation in Equations (1) through (3) and (6) through (9), then the number of salespersons is given and no longer a random variable. The relationship between optimality and post-optimality partitioning can also be viewed in the following way. Recall random-variable k is bounded between k^* and K in step 4, where k^* is the fixed number of aircraft to be flown in a priori analysis and K is the available fleet size. If $k = 1$, we have the PTSFLP; and if $k > 1$, we have the PMTSFLP. For the purpose of constructing *master* tours, an individual salesman-tour will be shorter as k increases, while the total tour-length among all salespersons will increase. In the limiting case when $k = n - 1$, or the number of salespersons is equal to the number of demand points, a spanning tree of $n - 1$ arcs is formed in which each salesperson visits one individual customer in an out-and-back fashion. Each salesperson spends the least amount of time on the road in this case but the company expends the largest travel budget to cover all salespersons. Aside from k, we recall the demand vector is another random variable for a priori analysis while it is fixed for a subset S of the I demand points in posteriori analysis. In steps 1 through 3 of the algorithm, k a-priori master-tours are constructed to cater for all demand-points. For post-optimization in step 4, we stay with the master-tour concept in which all potential demands in I are to be covered (irrespective of realization), and as k increases in step 4 of the PMTSFLP, the individual-salesman tour will for sure be shorter. Where the range constraint may be violated in step 3, the addition of salespersons on-the-road in step 4 will eventually afford a feasible solution provided that K is large enough and no single demand is beyond the range of an aircraft from home base.

Because the multiple-SFC is nothing more than the regular SFC with special partitioning, the asymptotic values of the regular SFC approximation will hold in both a priori and posteriori PMTSFLP. The reader is reminded that the order of visit in the "master" tour T remains the same irrespective of the partitioning. Also, we should think about the correspondence be-

tween the range of f and X–Y space in terms of proximity measures. Thus in the asymptotic case when there are ubiquitous demands over the X–Y space (which mapped into uniform distribution on the f-space), the partition among salesmen (and hence the value of k) becomes unimportant as far as solution accuracy is concerned. To see the implications of this result, consider two special cases. Case 1 is when the n demand points ($n \to \infty$) are randomly distributed on the Sierpinski curve. In this case, the clustering partition scheme gives the same result as the scheme where the curve is equally divided into k parts. The end product is $r = 1.25$. Consider case 2 when the n demand points are densely packed into k different clusters. In the limiting case, each cluster degenerates into a single point. Now the MSFC heuristic yields k "out-and-back" tours from the home depot. Then $r = 1$, or the clustering algorithm is perfect since no mis-classification is introduced due to distance distortions within a cluster during the SFC-transformation $f(X_i, Y_i)$.

V. CASE STUDY

The navigational-aid inspections by the Air Force to military facilities offer a unique opportunity to study the PMTSFLP. At the time of the original study (Merrill 1989), the Facilities Checking Squadron at Scott Air Force Base was scheduled to be operating six new C-29 aircraft to replace the then existing fleet. This created an environment for re-examining basing decisions. Operating cost and timeliness-in-response were only two among a myriad of considerations (which include political implications of base closings). But they will be the main concerns of this study. This in turn leads to the implicit assumption of considering only existing bases—rather than new bases—that can support such a fleet of inspection aircraft without additional capital investment. Four bases at the central part of the continental United States and two at the East and West coasts are candidates for close examination.

Historical data of flight-inspection missions over a one-year period were examined. The data were grouped by mission type and time frame. Only demands accomplished within ten duty-days of each other were grouped together considering crew-duty requirements. This is a reflection of timeliness-in-response as well as the common practice wherein a crew is out on-the-road for no more than ten days. Pure-inspection tours, operational-evaluation tours and the like were grouped separately to ensure a qualified crew could accomplish the tasking. Table I shows a reduced database where each mission contains no more than seven nodes. These were extracted from 1 September 1987 to 31 August 1988, forming a record for validation purposes using the analytical LP model (which can take only up to seven nodes). The bases shown in Table I are to be visited by one or more aircraft as specified.

All together there are 13 instances S (out of 28) consisting of subsets of bases requiring a visit by an inspection aircraft. Even though not shown in the data-set in Table I, it should be noted that Scott, the existing base, requires 15 visits during the year of study. In general, a zero frequency is assigned to the home base in the multiple-TSP analytical model (or its inspection needs are assumed to have been "taken care of,") whether it be Scott or other candidate home-base locations. However, the actual frequencies do play a role in locating median and center, and in employing the exact analytical MIP-model. While there are bases that have more frequent demands than others, the demand frequencies appear to be quite uniform in general. In other words, the p_i's are fairly close to be equal to one another on the average, with only minor exceptions.

An examination of their geographic locations also shows that the bases cover a majority of the continental United States. There appears to be no particular region which has an overwhelming demand in comparison to others. Similar to others in the 98 total, each base in the continental United States is identified by its longitude and latitude, from which great-circle distances between bases—following curvature of the globe—are readily computed. This paves the way for the implementation of SFC-heuristic through the use of HyperCard—a graphical package on the Apple MacIntosh computer.

While this constitutes a natural case study for the PMTSFLP, there are some restrictions that prevent a direct field implementation of the results reported here. Among the most prominent is the fact that these instances, or grouping of bases to be visited, are somewhat conditioned upon Scott being the base depot. But in our study here, these instances are used as demand for all alternate basing decisions.

This may be a good assumption for depots close to Scott in the central United States. But the assumption may become indefensible should a coastal base become a reality. Another problem is that one year of data may not be sufficient to cover all instances to validate *a priori* optimization as discussed here in this paper. As a result of all of these factors, the re-optimization algorithm will give us the theoretical tour length $G_n(j, T_k)$ in Section III, but not necessarily the actual distance and depot location as defined by real-world evaluation.

VI. RESULTS

As mentioned previously, space-filling-curve (SFC) heuristics were used to tackle the combined location-routing problem on potential demands that may exist in the 98 bases within the continental United States. We also extended the SFC heuristics to handle multiple-salesmen situations, implementing a mere suggestion by Bartholdi and Platzman (1988). In validating the multiple-SFC (MSFC) heuristics, results from the four candidate depots in the central part of the United States were compared with the optimal MTSP results based on the data in Table I. In 8 of the 13 instances for McConnell AFB and Altus AFB, 9 of 13 instances for Scott AFB, and 10 of 13 for Little Rock AFB, the results were exactly the same. For those cases where differences exist, the MSFC heuristics always provided answers within 17 percent of the optimal and averaged within 7.6 percent of the optimal between the four basing-decisions.

A. Route-Distance Savings. For the MTSP analytical-solutions, we investigated only 13 out of the 24 missions (instances) in the continental United States, in which six or fewer nodes are to be visited from the home base (see Table I). The 11 missions not analyzed include anywhere from 7 demands (served by 3 aircraft) to 19 demands (served by 4 aircraft), with a majority in the 9- to 10-demand category. The data-base analyzed represents 54 percent of the entire set of missions during the year. Limiting the data-base to six demands or less does not seem to bias the data geographically toward any portion of the continental United States. These are the missions, as explained above, in which optimal MTSP solutions are available, using the simplified set of subtour-breaking constraints contained in Equations (4)–(9). Bases chosen for basing considerations include six locations—four at the center part of the continental United States (Little Rock, McConnell, Altus and Scott), two at the coasts (Dover and Travis). The coastal locations are considered since in the larger single and multi-year data-base, occasional trips are made to Asia and Europe.

The results of the MSFC heuristics on Scott, the existing base of operation, indicated a distance savings of 2869 miles of en-route flying per mission over a year, assuming the same number of missions, the same number of aircraft dispatched per mission, and the same visiting locations. This translates to about three additional missions a year or seven additional flight-hours per mission. While this, by itself, is not a significant saving, the effect of *combined* location-and-routing changes can in fact be substantial. Figure 3 shows the various basing decisions. The existing basing-location at Scott is characterized by the horizontal line at an average of 6996 miles per tour under existing operating conditions. This compares with the various locational decisions at Little Rock, McConnell, Altus, Dover and Travis. It can be seen that the savings from improved routing at 2869 miles at Scott is significant enough to over shadow locational changes to the other bases except when the depot is located at Travis on the west coast. In reading Figure 3, one should keep in mind—once again—that the missions are grouped the same as in the mission data as shown in Table I, irrespective of the basing decision. As suggested in Table I, the number of aircraft used ranged from two to four. Should we allow the stops to interchange among missions, there would be substantially more savings involved, as alluded to above.

B. Consistency Among Base Locations. Table II compares the results of the combined location-routing analysis above with the median and center problems for the six candidate bases under consideration. Line two of the Table shows the exhaustive enumeration of the 13 instances of the candidate basing-locations. This means running the LP multiple-travelling-salesmen problem (MTSP) 78 times (Moore et

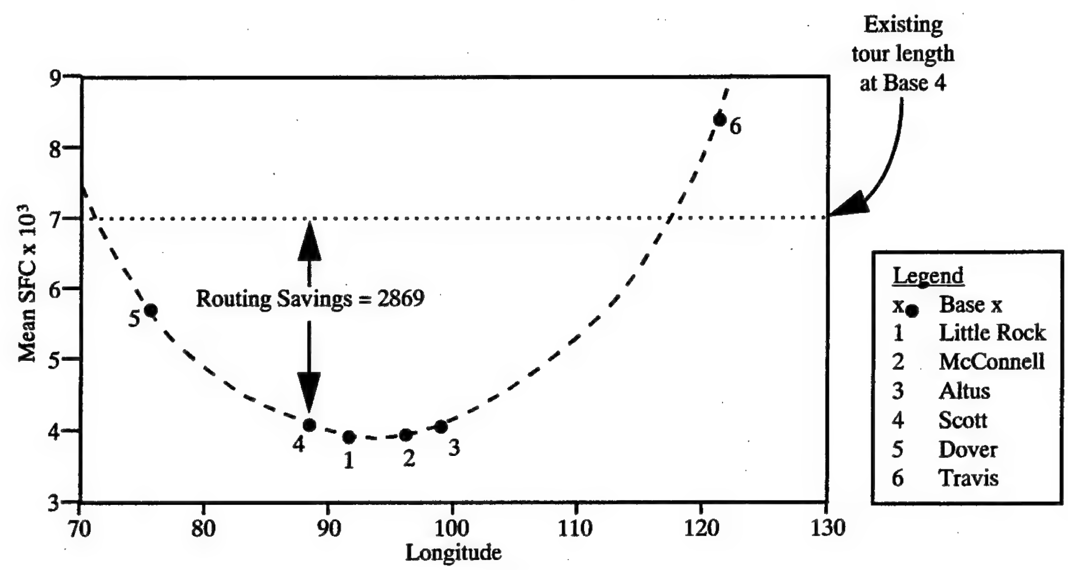

Figure 3 - "Comparison between existing basing location at Scott and alternate bases"

Table II. Base Ranking Using Location vs. Location-Routing Analyses

Base	Little Rock	McConnell	Altus	Scott	Dover	Travis
Location-Routing (MSFC)[2]	1	2	3	4	5	6
Location-Routing (MTSP)[3]	1	2	4	3	5	6
Median[4]	1	3	2	4	5	6
Center[4]	1	2	4	3	5	6

[2] The multiple-spacefilling-curve heuristic was run on the 98 bases within the continental United States.
[3] This analytical model was run on the 13 instances.
[4] The ranking was based on the 13-instance data-set.

al. 1991, Stephens et al. 1991, Shirley et al. 1992). The first line is the PMTSFLP as formulated above and solved by the MSFC heuristics. The third line is the median problem on the bases in the 13 instances. The fourth and last line is the center problem on the same data-set. It is interesting to note that the ranking of the six candidate bases are the same based both on the Heuristic MSFC and optimal MTSP solutions, except for the third and fourth, which are reversed between the two solution methods (Moore et al. 1991, Stephens et al. 1991). It is surprising to see the degree of similarity between the preference ranking among all four analyses. This is even more surprising considering that the "valley" appears to be quite "flat' in Figure 3, suggesting there is little difference between bases 1, 2, 3 and 4 at the central part of the United States. Perhaps the similarity between the median (for the 13 instances) and location-routing solutions can be explained in terms of an "evenly distributed" demand pattern. Likewise, the similarities between the center (for the 13 instances) and location-routing results can be attributed to a fairly symmetrical demand pattern about the candidate bases considered. Perhaps one can use the median and center as the approximation for locational deci-

sions in PMTSFLP when the demands are evenly distributed and are about equal. One can show that in the case of $p_i = p'$ and $n \to \infty$, 1-center, 1-median and PMTSFLP all approach and are all tightly distributed near the same depot location at the center of a unit square (Bartholdi and Platzman 1982).

Here a more formal justification is offered for the agreement between the median/center results and the PMTSFLP results. Recognize that the 1-median problem in the case of uniform demands reduces to $\min_{j \in I} \Sigma_{i \in I} d_{ij}$ and the 1-center problem reduces to $\min_{j \in I} (\max_{i \in I} d_{ij})$, which is the circumcenter. Both the median and center are located at the "dead center" of the unit square in this case. For PTSPs of uniform p, the length of the SFC tours is no greater than $2\sqrt{n}$ as in TSP, and can be considered to be $\sqrt{n}$ within a scaling constant for large n and MTSP. The symmetry of the problem also yields the central point of the unit square as the depot location for PMTSFLP. We have to contend that for the PMTSFLP (with $n \to \infty$), a near optimal solution is obtained by the following means. Consider an arbitrary depot, include its $k - 1$ nearest neighbors as $k - 1$ of the tours in T_k, and take the remaining $n - k$ points as the last tour in T_k. These nearest neighbors are all very close (within distance $O(\sqrt{k/n})$ of the depot) and the remaining large-tour solution is an instance of PTSP. Clearly, this yields a PMTSFLP location other than the dead center of a unit square. For this reason, only an asymptotic (rather than exact) statement can be made regarding the center, median and PMTSFLP being tightly distributed near the dead center of the square.

C. Computational Implications. The database analyzed in Table I represents 47 bases ($n = 47$) in the total of 98 included in the entire continental U.S. Equation (17) suggests that the accuracy of the MSFC is $O(\log n)$, the sample data is expected to yield an estimate

$$\frac{G_{47}(j^*, T_k^*)/w(j^{**})}{G_{98}(j^{**}, T_k^{**})/w(j^{**})} = \frac{O(\log 47)}{O(\log 98)}.$$

Inasmuch as there are at most 98 bases in the entire data-set, this ratio amounts to an accuracy of (log 47)/(log 98) = 84% of the best we can do. Once again, the amount of agreement between the locational ranking as obtained by MSFC and the remaining three analysis—MTSP, median and center—is striking. It appears that the only difference between MSFC, MTSP and the center is the rank reversal between Altus and Scott, two locations which are very close to one another in tour length according to Figure 3. The reversal between McConnell and Altus in the median ranking also reflects the minor difference between their tour lengths as shown again in Figure 3. Aside from the purpose of pre-positioning a home base, the MSFC analysis on the entire 98 bases can also be stored as in a "Roledex" file *a priori* for routing the inspection fleet on a day-to-day basis without any extra computation—a practice totally compatible with the field environment under which the squadron level is operating.

The probabilistic aspect of the PMTSFLP model is of such importance that we like to reemphasize the main thrust of this paper one more time. Line one of Table II represents the *a priori* tours constructed for real-world considerations when deterministic demands are never available. In such state of ignorance, the decision maker has no choice but to construct *a priori* tours out of the comprehensive set of bases which require flight checks—all 98 of them within continental United States. For validation purposes, we also show in line two *posterior* tours based on historical information on 13 instances of demands. It represents the tours selected *after* the demand realization been observed. The purpose is simply to show that the SFC heuristic can approximate an analytic location-routing model. It also serves to provide a necessary variant to justify the main point of this paper, which once again, is the construction of *a priori* tours in a world of uncertainty. Notice that line one is constructed from a random demand, but the number of aircraft (salespersons) is fixed. While the aircraft fleet is fixed in both line one and line two, it can be treated as a random variable as well, as illustrated in the PMTSFP algorithm and in the mathematical programs.

A comparison among all four lines in Table II is worthwhile, nevertheless the results need to be interpreted in terms of the data upon which the findings are based. Once again, the

data in Table I are historical records based upon Scott AFB being the home depot for all flight-inspection operations. But the experiments conducted in this research verify adequately how the PMTSFLP can be used to eventually base the fleet of inspection aircraft based on a much larger and more realistic validation data-base. It is understandable why a much larger database for line one does not deter the solution efficiency in that the MSFC, being of $O(n^2)$ complexity, is inherently a much simpler problem than the location-routing MTSP problem. The important point is that an extremely difficult problem can be approximated by a much easier problem.

D. Methodological Implications. A methodological finding is also worth discussing here. The simplified LP formulation for MTSP reported in Section II yielded solutions satisfactorily. In spite of the advances over the past decade in improving the compactness of the subtour-breaking constraint such as the one used in Equation (11), this formulation can be very practical for enumerative computation of all instances as performed in the MTSP solution above. It allows an integer solution to the MTSP using a regular linear-programming code. Remember that the number of times we solve the NP-hard MTSP is in the order of $n|S|$, where the number of instances $|S|$ is a large number. Any savings effort will be necessary. For the combined location-routing MIP as shown in the forcing/linking Equations (13) through (16), the number of additional equations beyond a multiple travelling-salesman-problem amounts to the order of n^2, which is a huge number as n becomes large.

Again, there is a difference between the two ways of solving the location-routing model. The MTSP way of enumerating all 13 instances among the alternative basing locations is the hard way, and the MSFC heuristic is the easy way. The former is NP-hard while the latter is $O(n^2)$. It is again gratifying to see that a difficult solution algorithm can be approximated by an easy algorithm. The efficiency of MSFC, when compared with the formal MTSP, can be dramatically brought out by the computer execution-times. Instead of several hours on the Apple Macintosh to run the 13 instances via the LP model, we are talking about 10 minutes using the MSFC heuristics to solve all 98 potential bases in the continental United States. Best of all, we need to do this only once, and the order of visit remains valid for all instances. It appears there is a definite role that SFC can play (Bertsimas 1989) in practical applications.

VII. CONCLUSION

The Probabilistic Travelling-Salesman Facility-Location Problem (PTSFLP) determines the whereabouts of a depot so as to minimize the tour among all demand nodes. We extended the PTSFLP to include multiple salesmen and offer a solution to the new model. We call this the Probabilistic Multiple-Travelling-Salesmen Facility-Location Problem (PMTSFLP). The deterministic multiple-travelling-salesmen-facility-location problem (MTSFLP) is formulated here for the first time, which helps to validate a particular instance of the PMTSFLP. This includes both a complete mixed-integer-program formulation incorporating range (distance) constraints and an integer-yielding linear-program formulation of the multiple-travelling-salesmen problem (MTSP) for small networks—the latter being used for complete enumeration of candidate depot locations in solving the MTSFLP. Through the use of Space-Filling Curves (SFC), an efficient multiple-SFC heuristic is constructed to solve the PMTSFLP. The inherent robustness of SFCs, particularly their ability to preserve the order of visit among all nodes, allows for *a priori* optimization of various instances of the PMTSFLP, where instances refer to situations where demands are realized.

The solution algorithm is by-and-large a straightforward mathematical extension of the single-travelling-salesman case using partitioning on the space-filling curve. However, it a significant extension in a couple of ways. First, the concept of random demands is extended to recognize a variable fleet-size (or a random number of salespersons). This provides flexibility for operators in the field to cope with mechanical breakdowns and crew-availability problems which often disable part of the fleet. Second, the multiple-salesmen case is the only meaningful location model to consider in our opinion. The locational problem disappears in the deterministic single-salesman case where the tour length does not depend on where the

base is located—thus reducing the problem simply to a routing problem.

Notice the fleet size is intimately related to the time *each* crew spends "on the road." The larger the number of aircraft used, the shorter each *individual-tour* of a crew. Conversely, the smaller the fleet, the longer each crew is away from home base in his/her round-robin tours. A small enough fleet-size may result in an unduly lengthy-tour for each aircraft (crew). It may even violate the "range constraint" in a vehicle-routing formulation of the PMTSFLP. To introduce feasibility, one has to increase the number of dispatched aircraft from k to $k+1$, $k+2, \ldots,$ up to the point when the time each crew spent on-the-road is brought down enough to fit the range constraint. Obviously, this is done at the expense of total system-mileage logged across the entire fleet. In a sentence, the individual-salesman tour is the shortest and the total mileage systemwide is the longest when $k=K$, where K is the maximally available fleet-size.

The asymptotic properties of the PMTSFLP algorithm, deemed an important finding of this paper, have been specified in terms of the performance bound $O(\log n)$, the accuracy bound 1.25, and the combined result of $r\,O(\log n)$, where n is a finite number of randomly-distributed demand-points and r ($1 \leq r \leq 1.25$) is problem-dependent constant. Another important contribution is that k, the actual number of salesmen deployed, is parametric in these performance bounds. An equivalent set of bounds can be written for a variable k (with $1 \leq k \leq K$), parametric on the probabilistic-demand vector $p=(p_1, p_2, \ldots, p_n)$.

The flight-inspection missions of the Air Force Facilities Checking Squadron lends itself to be an extremely natural case study of the PMTSFLP, particularly when the basing decision is re-examined with the acquisition of a new fleet of inspection aircraft. The case study is facilitated by the availability of one-year of historical data, covering a variety of instances with varying demand patterns. A datum for comparison also exists, consisting of the existing practice of the Squadron at Scott Air Force Base. Even though the one-year data and the way they are grouped does not allow coverage of all instances of demand to reach a definitive basing and routing decision, it does allow us to illustrate with clarity the efficiency of the multiple-SFC (MSFC) heuristic. For comparison, we solved the equivalent MTSP linear program using a simplified formulation of the subtour-breaking constraints. Of the 52 instances solved, 35 instances yielded the exact same solution as the MSFC heuristic. Overall, the heuristic solution is within 7.6 percent of the optimal for each basing-location on the average.

Our analysis shows that the PMTSFLP yields a ranking of depot locations which are very similar to the medians and centers. We have attributed this in part to the even-distribution of demand and the ubiquitous geographic-coverage of the demand nodes (Campbell 1992). Further analysis suggests there may in fact be some relationship between PMTSFLP and the center and median problems, perhaps in an asymptotic sense. This could allow the PMTSFLP depot-location to be approximated very easily by solving the 1-center or 1-median problem.

We have verified the accuracy of the MSFC heuristic through solving 78 multiple-TSPs among six alternate basing decisions. From an applicational standpoint, the research offers a fast, yet accurate analysis tool for basing and routing decisions. Given sufficient data, definite basing recommendations can be forthcoming. For the Facilities Checking Squadron, the SFC heuristic constitutes a convenient, yet rigorous, routing tool for the commander to respond to different demands on a day-to-day basis. Best of all, this can be carried out with minimal computational requirement. As such, it is a timely contribution toward current emphasis on quick-response analysis in the Department of Defense.

Even though this paper is motivated by defense applications, it is obvious that the methodology advanced here has equal applicability to many other time-sensitive delivery problems. These range from overnight-package delivery to emergency and disaster management when demands are random. All these problems require pre-positioning a depot and pre-planning vehicle routes that are robust enough for any contingencies that may arise.

ACKNOWLEDGMENT

Special thanks to all the colleagues in Scott Air Force Base that rendered invaluable assistance during the conduct of this study, including (in alphabetical order) C. Austin, M. Ball, M. Donnelly, E. Miller, M. Pruden, and J. Sousaris. The authors also would like to ac-

knowledge the ideas offered by A. Odoni and T. F. Schuppe at the formative stages of this research. Our appreciation also goes to the 1991 Koopman Prize Award Committee of the Military Applications Section of the Operations Research Society of America (now the Institute of Operations Research and Management Science). They gave us the heart-warming encouragement by their conferral of the prestigious award on the first draft of this paper. As we refine the research paper for archival journal publication, the authors benefited greatly from the insightful and constructive suggestions of the referees, which significantly improve the quality of this paper. Obviously, the statements made in this paper are strictly those of the authors. They do not necessarily represent those of their present and past affiliations.

REFERENCES

Bartholdi, J. and Platzman, L. (1988). "Heuristics based on spacefilling curves for combinatorial problems in Euclidean space", *Management Science*, Vol. 34, No. 3, pp. 291–305.

Bartholdi, J., Platzman, L., Collins, R., and Warden, W. (1983). "A minimal technology routing system for meals on wheels", *Interfaces*, Vol. 13, pp. 1–8.

Bartholdi, J. J., and Platzman, L. K. (1982). "An O(n log n) planar travelling salesman heuristic based on space filling curves", *Operations Research Letters*, Vol. 1, pp. 121–125.

Berman, O., and Simchi-Levi, D. (1988). "A Heuristic Algorithm for the Travelling Salesman Location Problem on Networks", *Operations Research*, Vol. 36, pp. 478–484.

Berman, O., and Simchi-Levi, D. (1988a). "Finding the Optimal *A Priori* Tour and Location of a Travelling Salesman with Non-Homogeneous Customers", *Transportation Science*, Vol. 22, pp. 148–154.

Bertsimas, D. J. (1988). "Probabilistic combinatorial optimization problems," Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, Mass.

Bertsimas, D. (1989). "Traveling Salesman Facility Location Problems," *Transportation Science*, Vol. 23, pp. 184–191.

Bertsimas, D. J. (1992). "A vehicle routing problem with stochastic demand," *Operations Research*, Vol. 40, No. 3, pp. 574–585.

Bertsimas, D., Chern, P., and Peterson, M. (1995). "Computational approaches to stochastic vehicle routing problems", *Transportation Science*, Vol. 29, pp. 342–352.

Bertsimas, D. J., and G. Van Ryzin (1991). "A stochastic and dynamic vehicle routing problem in the Euclidean plane", *Operations Research*, Vol. 39, pp. 601–615.

Bertsimas, D. J., Jaillet, P., and Odoni, A. R. (1990). "A priori optimization", *Operations Research*, Vol. 38, pp. 1019–1033.

Bertsimas, D. J., and Simchi-Levi, D. (1996). "A new generation of vehicle routing research: robust algorithms, addressing uncertainty", *Operations Research*, Vol. 44, pp. 286–304.

Bowerman, R. L., and Calamai, P. H. (1994). "The spacefilling curve with optimal partitioning heuristic for the vehicle routing problem", *European Journal of Operational Research*, Vol. 76, pp. 128–142.

Bramel, J. (1992). "The capacitated vehicle routing problem with unsplit demands: Practical algorithms and analytical results", working paper, Department of Industrial Engineering and Operations Research, Columbia University, New York, New York.

Bramel, J., Coffman Jr., E. G., Shor, P. W., and Simchi-Levi, D. (1992). "Probabilistic analysis of the capacitated vehicle routing problem with unsplit demands", *Operations Research*, Vol. 40, No 6, pp. 1095–1106.

Campbell, J. F. (1992). "Continuous and discrete demand hub location problem", working paper CBIS 92-06-01, Center for Business and Industrial Studies, University of Missouri, St. Louis, Missouri.

Chan, Y., and Baker, S. (1996). "The multiple depot multiple travelling salesmen problem: vehicle range and servicing frequency implementations" working paper,

Department of Operational Sciences. Air Force Institute of Technology, Wright-Patterson AFB, Ohio.

Chan, Y. (forthcoming). *Facility Location and Land Use: Multicriteria Analysis of Spatial Temporal Information*, ITP/Southwestern.

Daskin, M. (1995). *Network and Discrete Location: Models, Algorithms and Applications*, Wiley Interscience.

Dror, M., Laporte, G., and Trudeau, P. (1989). "Vehicle routing with stochastic demands: properties and solution frameworks", *Transportation Science*, Vol. 23, pp. 166–176.

Hakimi, S. (1964). "Optimal location of switching centers and the absolute center and the medians of a graph", *Operations Research*, Vol. 12, pp. 450–459.

Hargrave, W., and Nemhauser, G. (1962). "On the Relation Between the Travelling Salesman Problem and the Longest Path Problem", *Operations Research*, Vol. 10, pp. 647–657.

Hoffman, A., and Wolfe, P. (1985). "History", in the *Travelling Salesman Problem*, edited by E. Lawler et al., Chicester, GB: John Wiley & Sons.

Jamil, M., Batta, R., and Malon, D. M. (1994). "The traveling repairperson home base location problem", *Transportation Science*, Vol. 28, pp. 150–160.

Laporte, G., Louveau, F. V., and Mercure, H. (1994). "A priori optimization of the probabilistic traveling salesman problem", *Operations Research*, Vol. 42, pp. 543–549.

Laporte, G., Louveau, F., and Mercure, H. (1987). "Models and Exact Solutions for a Class of Stochastic Location-Routing Problems", G-87-14, Ecole des Hautes Etudes Commerciale, University of Montreal, Montreal, Canada.

Loftus, M., and Chan, Y. (1992). "A combined location and routing model", working paper, Department of Operational Sciences, Air Force Institute of Technology, Wright-Patterson AFB, Ohio.

Merrill, D. (1989). "Facility location and routing to minimize the en route distance of flight inspection missions", master's thesis submitted to the Department of Operational Sciences, Air Force Institute of Technology, Wright-Patterson AFB, Ohio.

Moore, R. S., Puhek, P., and Chan, Y. (1991). "A travelling salesman/facility location problem", working paper, Dept. of Operational Sciences, Air Force Institute of Technology, Wright-Patterson AFB, Ohio.

Reynolds, J. K., Baker, S. B., and Chan, Y. (1990). "Multiple travelling salesman problem: integer programming, subtour breaking and time constraints", working paper for Oper 6.63, Dept. of Operational Sciences, Air Force Institute of Technology, Wright-Patterson AFB, Ohio.

Shirley, M., and Chan, Y. (1993). "Single run facility location study", working paper, Department of Operational Sciences, Air Force Institute of Technology, Wright-Patterson AFB, Ohio.

Shirley, M., Sosebee, B., and Chan, Y. (1992). "Facility location study", working paper, Department of Operational Sciences, Air Force Institute of Technology, Wright-Patterson AFB, Ohio.

Stephens, A., Habash, N., and Chan, Y. (1991). "Alternate depot: Altus Air Force Base", working paper, Department of Operational Sciences, Air Force Institute of Technology, Wright-Patterson AFB, Ohio.

Call for Presentations

1st MAS National Meeting

19-21 May 1997

Huntsville Alabama

Going Forward Into The Future With Military Operations Research

Sessions and presentations in the area of contemporary military operations research and management science are solicited. Prospective participants and attendees may contact either the General Chair, Anthony Brinkley, Teledyne Brown Engineering, Huntsville, AL (tony_brinkley@pobox.tbe.com or (205) 726-5857), or the Program Chair , Bruce Fowler, US Army Aviation Missile Command, ATTN: AMSAM-RD-AS, Redstone Arsenal, AL 35898 (fowler-bw@redstone.army.mil or (205) 876-8173.)

Hosted by the Redstone Arsenal-Huntsville Military Operations Research Section of MAS

Application and Extension of the Thruput II Optimization Model for Airlift Mobility

Richard E. Rosenthal
Steven F. Baker
Lim Teo Weng
David F. Fuller
David Goggins
Ayhan O. Toy
Yasin Turker

Operations Research Department
Naval Postgraduate School
1411 Cunningham Rd.
Monterey, California 93943-5219

David Horton
Daniel Briand

U.S. Air Force Studies and Analyses Agency

David P. Morton

Graduate Program in Operations Research
University of Texas at Austin

APPLICATION AREA:
Mobility & Transport of Forces

OR METHODOLOGY:
Linear Programming

ABSTRACT

THRUPUT II is a linear programming model developed at the Naval Postgraduate School for the U.S. Air Force Studies and Analyses Agency (AFSAA) to help improve the efficiency of the airlift mobility system. It determines the maximum on-time throughput of cargo and passengers that can be transported with a given aircraft fleet over a given network, subject to appropriate physical and policy constraints. THRUPUT II was used in the analysis provided by AFSAA to the C-17 Defense Acquisition Board in November, 1995. This paper reviews the model's formulation, describes its use in the C-17 analysis, and reports extensions that have been developed since the model's first appearance.

1 INTRODUCTION

This paper is a status report on a multiyear research effort to apply optimization modeling technology to the analysis of strategic airlift mobility. The purpose of the research is to help the U.S. Air Force improve logistical efficiency. Optimization is used to determine the maximum on-time throughput of cargo and passengers that can be transported with a given aircraft fleet over a given network, subject to appropriate physical and policy constraints. The model can be used to help answer questions about selecting airlift assets and about investing or divesting in airfield infrastructure.

The primary model discussed in this paper is called THRUPUT II, which was introduced in a Naval Postgraduate School (NPS) Masters thesis [Lim, 1994] and further developed in a *Military Operations Research* article [Morton, Rosenthal, and Lim, 1996]. Since those earlier publications were written, THRUPUT II provided inputs to the C-17 Defense Acquisition Board decision of November 1995. This experience and other subsequent developments are covered here. A new model is currently under joint development between NPS and the RAND Corporation [Melody *et al*, 1996]. The distinguishing features of this new model are discussed in the conclusion.

The progenitors of THRUPUT II were the first THRUPUT, developed at the Air Force Studies and Analyses Agency [Yost, 1994]; and the Mobility Optimization Model (MOM), developed at the Joint Staff's Force Structure Resource, and Assessment Directorate (J8) [Wing *et al*, 1991]. All of these models are implemented with the General Algebraic Modeling System (GAMS) [Brooke, Kendrick and Meeraus, 1992].

Examples of the types of mobility questions that can be analyzed with optimization are: For a given fleet and a given network,

- Are the aircraft and airfield assets adequate for the deployment scenario?
- What are the impacts of shortfalls in airlift capability?
- Where are the system bottlenecks and when will they become noticeable?

This type of analysis can be used to help answer questions about selecting airlift assets and about investing or divesting in airfield infrastructure. Such analyses are accomplished through repeated runs of the model. Each run assumes a particular scenario as defined by a given set of time-phased movement requirements and a given set of available aircraft and airfield assets. It is then solved for optimal values for the number of missions flown, and the amounts of cargo and passengers carried, for each unit, by each aircraft type, via each route, in each time period.

After describing the optimization model in Sections 2 and 3, Sections 4 and 5 discuss analyses performed in the recent non-developmental airlift aircraft (NDAA)/C-17 study. A special-purpose algorithm for solving large problem instances and a modeling extension to incorporate aircraft reliability are described in Section 6. Section 7 presents conclusions and ongoing research.

2 OVERVIEW OF MODEL

In this section we give a conceptual overview of the airlift optimization model. Then, Section 3 provides a detailed mathematical formulation. Sections 2 and 3 can be skipped by readers familiar with [Morton, Rosenthal, and Lim, 1996].

2.1 Model Features

The model has been designed to handle many of the airlift system's particular features and modes of operation. For example, the payload an aircraft can carry depends on the maximum leg distance of a mission (shorter mission legs allow greater payloads), and aircraft with heavy loads may

be required to make frequent enroute stops. Also, there is a need to ensure cargo-to-carrier compatibility since some military hardware is too bulky to fit into certain aircraft. These features have been incorporated in the model to make it as realistic as possible. Others, such as the use of tanker aircraft for aerial refueling of airlift aircraft, incorporating crews, and modeling intra-theater shuttles and ground transportation are the subject of the follow-on model mentioned earlier. The major features of the airlift system currently captured by the model include:

- Multiple origins and destinations: In contrast to MOM, the current model allows the airlift to use multiple origin, enroute and destination airfields.
- Flexible routing structure: The air route structure supported by the model includes delivery and recovery routes with a variable number of enroute stops (usually between zero and three). This provision allows for a mixture of short-range and long-range aircraft. The model can thus analyze trade-offs between higher-payload, shorter-range flights and lower-payload, longer-range flights. For further routing flexibility, the model also allows the same aircraft to fly different delivery and recovery routes on opposite ends of the same mission.
- Aircraft-to-route restrictions: The user may impose aircraft-to-route restrictions; e.g., only military aircraft may use military airfields for enroute stops. This particular provision arises because the USAF Air Mobility Command (AMC) may call upon civilian commercial airliners to augment USAF aircraft in a deployment, under the Civil Reserve Airfleet (CRAF) program. The model distinguishes between USAF and CRAF aircraft.
- Aircraft assets can be added over time: This adds realism to the model, because CRAF and other aircraft may take time to mobilize and are typically unavailable at the start of a deployment.
- Delivery time windows: In a deployment, a unit is ready to move on its available-to-load date (ALD) and has to arrive in the theater by its required-delivery-date (RDD). This aspect of the problem has been incorporated in the model through user-specified time windows for each unit. The model treats the time windows as "elastic" in that cargo may be delivered late, subject to a penalty.

2.2 Conceptual Model Formulation

The primary decision variables are the number of missions flown, and the amounts of cargo and passengers carried, for each unit, by each aircraft type, via each available route, in each time period. Additional variables are defined for the recovery flights, for aircraft inventoried at airfields, and for the possibility (at high penalty cost) of not delivering required cargos or passengers.

2.2.1 Objective Function

The purpose of the optimization model is to maximize the effectiveness of the given airlift assets, subject to appropriate physical and policy constraints. The measure of effectiveness is the minimization of total weighted penalties incurred for late deliveries and non-deliveries. The penalties are weighted according to two factors: the priority of the unit whose movement requirement is not delivered on time, and the degree of lateness. The penalty increases with the amount of time late, and non-delivery has the most austere penalty.

The anticipated use of the model is for situations when the given airlift resources are insufficient for making all the required deliveries on time. On the other hand, if there are enough resources for complete on-time delivery, then the model's secondary objective function is to choose a feasible solution that maximizes unused aircraft. The motivation of the secondary objective is that if the available aircraft are used as frugally as possible, while still meeting the known demands and observing the known constraints, then the mobility system will be as well prepared as it can be for unplanned breakdowns and unforeseen requirements, such as additional contingencies.

2.2.2 Constraints

The model's constraints can be grouped into the five categories: demand satisfaction, aircraft balance, aircraft capacity, aircraft utilization, and aircraft handling capacity at airfields.

- Demand Satisfaction Constraints: The cargo demand constraints attempt to ensure for each unit that the correct amounts of cargo move to the required destination within the

specified time window. The passenger demand constraints do the same for each unit's personnel. The demand constraints have elastic variables for late delivery and non-delivery. The optimization will seek to avoid lateness and non-deliveries if it is possible with the available assets, or to minimize them if not.

- Aircraft Balance Constraints: These constraints keep physical count of aircraft by type (e.g., C-17, C-5, C-141, etc.) in each time period. They ensure that the aircraft assets are used only when they are available.
- Aircraft Capacity Constraints: There are three different kinds of constraints on the physical limitations of aircraft—troop carriage capacity, maximum payload, and cabin floor space—which must be observed at all times.
- Aircraft Utilization Constraints: These constraints ensure that the average flying hours consumed per aircraft per day are within AMC's established utilization rates for each aircraft type.
- Aircraft Handling Capacity at Airfields: These constraints ensure that the number of aircraft routed through each airfield each day is within the airfield's handling capacity.

2.3 Assumptions

Some major assumptions of the model are listed below. These are known to be sacrifices of realism, but such assumptions are needed in modeling most real-world problems due to the limitations of data availability or the need to avoid computational intractability.

- Air Force planners use a measure called Maximum-on-Ground (MOG) to represent airfield capacity. The literal translation of MOG as the maximum number of planes that can be simultaneously on the ground at an airfield is somewhat misleading, because the term MOG means more than just the number of parking spaces at an airfield. In actuality, airfield capacity depends on many dimensions in addition to parking, including material handling equipment, ground services capacity and fuel availability. Some Air Force planners use the terms parking MOG and working MOG to distinguish between parking space limits and servicing capability. Working MOG is always smaller than parking MOG, and is the only MOG for which we have data. Working MOG is an approximate measure because it attempts to aggregate the capacities of several kinds of services into a single, unidimensional figure. Disaggregation of airfield capacity into separate capacities for parking spaces and for each of the specific services available would yield a more accurate model. Ongoing projects at AMC [Schubert, Whisman, and Steppe, 1996] and RAND [Stucker, 1996] involve stochastic and deterministic simulations, respectively, whose purpose is to determine appropriate, and possibly multidimensional, MOG values. The model presented here will benefit from these investigations.
- Inventoried aircraft at origin and destination airfields are considered not to affect the aircraft handling capacity of the airfield. This assumption is not strictly valid since an inventoried aircraft takes up parking space, but, as noted, working MOG dominates parking MOG.
- Deterministic ground time: Aircraft turnaround times for onloading and offloading cargo and enroute refueling are assumed to be known constants, although they are naturally stochastic. This ignores the fact that deviations from the given service time can cause congestion on the ground. To offset the optimism of this assumption, an efficiency factor is used in the formulation of aircraft handling capacity constraints to cushion the impact of randomness. Then, in Section 6, we describe a stochastic optimization formulation that explicitly models stochastic ground times and indicate how this optimization model has been linked with a discrete-event simulation.

Other approximations of reality employed in the model for computational tractability are aggregation of airfields, discretization of time, and continuous decision variables. A limitation on the scope of the model is that it considers only inter-theater, not intra-theater deliveries.

3 OPTIMIZATION MODEL

This section gives a mathematical formulation of the conceptual model outlined above. The airlift optimization model is formulated as a multi-period, multi-commodity network-based linear program (LP) with a large number of side constraints. Two key concepts are employed in the model. The first is the use of a

time index to track the locations of aircraft for each time period. The modeling advantages of knowing when an aircraft will arrive at a particular airfield are that it enables us to model aircraft handling capacity at airfields and to determine unit closure (i.e., the time when all of a unit's deliveries are complete). This approach is in contrast to the THRUPUT model of [Yost, 1994], which takes a static-equilibrium or steady-state approach.

The second key concept is model reduction through data aggregation and the removal of unnecessary decision variables and constraints prior to optimization. This is necessary as the airlift problem is potentially very large. Without this model reduction step, the number of decision variables would run into the millions, even for a nominal deployment. The unnecessary decision variables and constraints are removed by extensive checking of logical conditions, performed by GAMS during model generation. This is discussed in greater detail in Section 5.

3.1 Indices

u	indexes units, e.g., 82nd Airborne
a	indexes aircraft types, e.g., C-17, C-5, C-141
t,t'	indexes time periods
b	indexes all airfields (origins, enroutes, and destinations)
i	indexes origin airfields
k	indexes destination airfields
r	indexes routes

3.2 Index Sets

Airfield Index Sets

B	set of available airfields
$I \subseteq B$	origin airfields
$K \subseteq B$	destination airfields

Aircraft Index Sets

A	set of available aircraft types
$A_{bulk} \subseteq A$	aircraft capable of hauling bulk-sized cargo
$A_{over} \subseteq A_{bulk}$	aircraft capable of hauling over-sized cargo
$A_{out} \subseteq A_{over}$	aircraft capable of hauling out-sized cargo

Bulk cargo is palletized on 88 × 108 inch platforms (84 × 104 usable) and can fit on any military aircraft (as well as cargo-configured CRAF) [Merrill, 1997]. Over-sized cargo is non-palletized rolling stock; it is larger than bulk cargo and can fit on a C-141, C-5 or C-17. Out-sized cargo is very large non-palletized cargo that can fit into a C-5 or C-17 but not a C-141.

Route Index Sets

R	set of available routes
$R_a \subseteq R$	permissible routes for aircraft type a
$R_{ab} \subseteq R_a$	permissible routes for aircraft type a that use airfield b
$R_{au} \subseteq R_a$	permissible routes for aircraft type a carrying cargo or troops for unit u
$R_{ai} \subseteq R_a$	permissible routes for aircraft type a that use origin airfield i
$R_{ak} \subseteq R_a$	permissible routes for aircraft type a that use destination airfield k
$DR_i \subseteq R$	delivery routes that originate from origin i
$RR_k \subseteq R$	recovery routes that originate from destination k

A *delivery route* is a route flown from a specific unit's origin to its destination for the purpose of delivering cargo and/or passengers. A *recovery route* is a route flown from a unit's destination to that unit's or some other unit's origin, for the purpose of making another delivery. Since recovery flights carry much less weight than deliveries, the recovery routes from k to i may have fewer enroute stops than the delivery routes from i to k.

Time Index Sets

T	set of time periods
$T_{uar} \subseteq T$	possible start times for aircraft of type a flying a mission for unit u on route r

The set T_{uar} covers the allowed time window for unit u, which starts on the unit's available-to-load date and ends on the unit's re-

quired delivery date, plus some extra time up to the maximum allowed lateness for the unit.

3.3 Given Data

Movement Requirements Data

$MovePAX_u$	Troop movement requirement for unit u
$MoveUE_u$	Equipment movement requirement in short tons (stons) for unit u
$ProBulk_u$	Proportion of unit u cargo that is bulk-sized
$ProOver_u$	Proportion of unit u cargo that is over-sized
$ProOut_u$	Proportion of unit u cargo that is out-sized

Penalty Data

$LatePenUE_u$	Lateness penalty (per ston per day) for unit u equipment
$LatePenPAX_u$	Lateness penalty (per soldier per day) for unit u troops
$NoGoPenUE_u$	Non-delivery penalty (per ston) for unit u equipment
$NoGoPenPAX_u$	Non-delivery penalty (per soldier) for unit u troops
$MaxLate$	Maximum allowed lateness (in days) for delivery
$Preserve_{at}$	Penalty (small artificial cost) for keeping aircraft type a in mobility system at time t

Cargo Data

$UESqFt_u$	Average cargo floor space (in sq. ft.) per ston of unit u equipment
$PAXWt$	Average weight of a soldier inclusive of personal equipment

Aircraft Data

$Supply_{at}$	Number of aircraft of type a that become available at time t
$MaxPAX_a$	Maximum troop carriage capacity of aircraft type a
$PAXSqFt_a$	Average cargo space (in sq. ft.) consumed by a soldier for aircraft type a
$ACSqFt_a$	Cargo floor space (in sq. ft.) of aircraft type a
$LoadEff_a$	Cargo space loading efficiency (<1) for aircraft type a. This accounts for the fact that it is not possible in practice to fully utilize the cargo space.
$URate_a$	Established utilization rate (flying hours per day) for an aircraft of type a

Airfield Data

$MOGCap_{bt}$	Aircraft capacity (in narrow-body equivalents) at airfield b in time t
$MOGReq_{ab}$	Conversion factor to a narrow-body equivalent for an aircraft of type a at airfield b
$MOGEff$	MOG efficiency factor (<1), to account for the fact that it is impossible to fully utilize available MOG capacity due to randomness of ground times

Aircraft Route Performance Data

$MaxLoad_{ar}$	Maximum payload (in stons) for aircraft type a flying route r
$GTime_{abr}$	Aircraft ground time (due to onload or offload of cargo, refueling, maintenance, etc.) needed for aircraft type a at airfield b on route r
$DTime_{abr}$	Cumulative time (flight time plus ground time) taken by aircraft type a to reach airfield b along route r
$FltTime_{ar}$	Total flying hours consumed by aircraft type a on route r
$Ctime_{ar}$	Cumulative time (flight time plus ground time) taken by aircraft type a on route r
$DaysLate_{uart}$	Number of days late unit u's requirement would be if delivered by aircraft type a via route r with mission start time t

3.4 Decision Variables

Mission Variables

X_{uart} Number of aircraft of type a that airlift unit u via route r with mission start time t

Y_{art} Number of aircraft of type a that recover from a destination airfield via route r with start time t

Aircraft Allocation and De-allocation Variables

$Allot_{ait}$ Number of aircraft of type a that are allocated to origin i at time t

$Release_{ait}$ Number of aircraft of type a that were allocated to origin i prior to time t but are not scheduled for any missions from time t on

Aircraft Inventory Variables

H_{ait} Number of aircraft of type a inventoried at origin i at time t

HP_{akt} Number of aircraft of type a inventoried at destination k at time t

$Nplanes_{at}$ Number of aircraft of type a in the air mobility system at time t

Airlift Quantity Variables

$TonsUE_{uart}$ Total stons of unit u equipment airlifted by aircraft of type a via route r with mission start time during period t

$TPAX_{uart}$ Total number of unit u troops airlifted by aircraft of type a via route r with mission start time during period t

Elastic (Nondelivery) Variables

$UENoGo_u$ Total stons of unit u equipment not delivered in the prescribed time frame

$PAXNoGo_u$ Number of unit u troops not delivered in the prescribed time frame

Each of the decision variables is constrained to be non-negative.

3.5 Formulation of the Objective Function

minimize:

$$\sum_u \sum_a \sum_{r \in R_a} \sum_{t \in T_{uar}} LatePenUE_u$$
$$\cdot DaysLate_{uart} \cdot TonsUE_{uart}$$
$$+\sum_u \sum_a \sum_{r \in R_a} \sum_{t \in T_{uar}} LatePenPAX_u$$
$$\cdot DaysLate_{uart} \cdot TPAX_{uart}$$
$$+\sum_u (NoGoPenUE_u \cdot UENoGo_u$$
$$+NoGoPenPAX_u \cdot PAXNoGo_u)$$
$$+\sum_a \sum_t Preserve_{at} \cdot NPlanes_{at}$$

The $DaysLate_{uart}$ penalty parameter has value zero if $t + CTime_{ar}$ is within the prescribed time window for unit u. Thus, the first two terms of the objective function take effect only when a delivery is late. The third term in the objective function corresponds to cargo and passengers that cannot be delivered even within the permitted lateness. Late delivery and non-delivery occur only when airlift assets are insufficient for on-time delivery.

The reason for including elastic variables that allow late delivery and non-delivery is to ensure that the model produces useful information even when the given assets are inadequate for the given movement requirements. The alternative of using an inelastic model (i.e., a model with hard constraints that insist upon complete on-time delivery) is inferior because it would report infeasibility without giving any insight about what can be done with the assets available.

A useful modeling excursion that is made possible by the elastic variables is to vary the number of time periods. As the horizon is shortened, it is interesting to observe the increase in lateness and non-delivery.

As noted, the model's anticipated use is in cases when the airlift assets are insufficient for full on-time delivery. In the opposite case, the model will be governed by the fourth term of

the objective function, which rewards asset preservation for the reasons given in Section 2.

Some care must be taken in selecting the lateness and non-delivery penalties and the aircraft preservation rewards to ensure consistency. Late delivery should be preferred to non-delivery. The weights will be consistent with this preference provided the late penalty (per ston per day) is less than the corresponding non-delivery penalty (per ston) divided by the maximum allowed lateness (in days).

3.6 Formulation of the Constraints

As noted in the conceptual model, there are five categories of constraints. Their mathematical formulations are as follows.

3.6.1 Demand Satisfaction Constraints

There are four different kinds of demand constraints, corresponding to troops and the three classes of cargo (bulk, over-sized and out-sized). Separate constraints are required for the different cargo types to ensure cargo-carrier compatibility. For example, a carrier of over-sized cargo cannot be used to carry the larger out-sized cargo. On the other hand, it is possible to use a carrier of out-sized cargo to carry over-sized cargo. The model accounts for this asymmetry.

The demand constraints also account for the desired delivery time-windows by use of the index sets T_{uar} and the lateness parameters $DaysLate_{uart}$.

3.6.2 Aircraft Balance Constraints

There are five kinds of aircraft balance constraints enforced for each aircraft type in each time period. At origin airfields, they ensure that the number of aircraft assigned for delivery missions plus those inventoried for later use plus those put in the released status equal the number inventoried from the previous period plus recoveries from earlier missions and the new supply of aircraft that is allocated to the origin.

The meaning of releasing, or de-allocating, an airplane in period t is that it is not flown on any missions from period t through the end of the horizon. In practice, the analyst can interpret a release in the model's solution in a variety of ways. It can mean, as in the case of the civilian CRAF aircraft, that the plane is literally sent back to its owner, but not necessarily. The aircraft can also be kept in the mobility system, available as a replacement in case of breakdowns or for unforeseen demands.

The second kind of aircraft balance constraints concerns destinations. They are similar to the first kind except releases are not allowed and the roles of delivery and recovery missions are reversed. The third kind of aircraft balance constraint ensures that if any new planes become available in period t, they are allotted appropriately among the origins. There is a potential gain in efficiency to allow the optimizer to make these allocation decisions, rather than relying on the user to pre-assign them to origin airfields. The fourth type of aircraft bal-

Demand Satisfaction Constraints for All Classes of Cargo:

$$\sum_{a\in A_{bulk}} \sum_{r\in R_{au}} \sum_{t\in T_{uar}} TonsUE_{uart} + UENoGo_u = MoveUE_u \quad \forall u \text{ with } MoveUE_u > 0$$

Demand Satisfaction Constraints for Out-Sized Cargo:

$$\sum_{a\in A_{out}} \sum_{r\in R_{au}} \sum_{t\in T_{uar}} TonsUE_{uart} + UENoGo_u \geq ProOut_u \cdot MoveUE_u \quad \forall u \text{ with } MoveUE_u > 0$$

Demand Satisfaction Constraints for Over-Sized Cargo:

$$\sum_{a\in A_{ovr}} \sum_{r\in R_{au}} \sum_{t\in T_{uar}} TonsUE_{uart} + UENoGo_u \geq (ProOver_u + ProOut_u) \cdot MoveUE_u \quad \forall u \text{ with } MoveUE_u > 0$$

Demand Satisfaction Constraints for Troops:

$$\sum_{a} \sum_{r\in R_{au}} \sum_{t\in T_{uar}} TPAX_{uart} + PAXNoGo_u = MovePAX_u \quad \forall u \text{ with } MovePAX_u > 0$$

ance constraints is a set of accounting equations for defining the $NPlanes_{at}$ variables based on cumulative allocations and releases.

In the following constraints we use the notation $[Ctime_{ar}]$ to denote $Ctime_{ar}$ rounded to the nearest integer.

Aircraft Balance Constraints at Origin Airfields:

$$\sum_{u}\sum_{r\in DR_i} X_{uart} + H_{ait} + Release_{ait}$$
$$= H_{ai,t-1} + Allot_{ait}$$
$$+ \sum_{r\in R_{ai}} \sum_{t'+[CTime_{ar}]=t} Y_{art'} \quad \forall a,i,t$$

Aircraft Balance Constraints at Destination Airfields:

$$\sum_{r\in RR_k} Y_{art} + HP_{akt} = HP_{ak,t-1}$$
$$+ \sum_{u}\sum_{r\in R_{ak}} \sum_{\substack{t'\in T_{uar}\\ t'+[CTime_{ar}]=t}} X_{uart'} \quad \forall a,k,t$$

Aircraft Balance Constraints for Allocations to Origin Airfields:

$$\sum_{t'=1}^{t}\sum_{i} Allot_{ait'} \le \sum_{t'=1}^{t} Supply_{at'} \quad \forall a,t$$

The above constraint is in the cumulative form, rather than in the simpler form $\Sigma i\ Allot_{ait} \le Supply_{at}$ to allow aircraft that become available in period t to be put into service at a later period.

Aircraft Balance Constraints Accounting for Allocations and Releases:

$$NPlanes_{at} = \sum_{t'=1}^{t}\sum_{i} Allot_{ait'}$$
$$- \sum_{t'=1}^{t}\sum_{i} Release_{ait'} \quad \forall a,t$$

The fifth and final set of aircraft balance constraints helps to correct the discretization error that can result from rounding $Ctime_{ar}$ to $[Ctime_{ar}]$, the nearest integer, in the other balance constraints. For example, suppose $Ctime_{ar}$ is less than half a day for some aircraft a and route r. When this time is rounded to zero in the balance constraints of the route's origin and destination, these constraints unrealistically permit an unlimited number of missions per day on that route. Solving the model with this deficiency would yield overly optimistic results.

One way to fix this problem would be to insist that $Ctime_{ar}$ be rounded up to a higher integer. Then the model would be overly pessimistic, because it would rule out the possibility of an aircraft flying two or more missions in a day even when this is possible. This sort of problem is common in mathematical modeling whenever time is discretized. The approach taken here is to enforce the following additional constraints, based on the cumulative plane-days available.

Cumulative Aircraft Balance Constraints:

$$\sum_{r\in R_a}\sum_{t'=1}^{t}\sum_{u} K_{artt'} X_{uart'}$$
$$+ \sum_{r\in R_a}\sum_{t'=1}^{t} K_{artt'} Y_{art'} + \sum_{i}\sum_{t'=1}^{t} H_{ait'}$$
$$+ \sum_{k}\sum_{t'=1}^{t} HP_{akt'} \le \sum_{t'=1}^{t} NPlanes_{at'} \quad \forall a,t$$

where $K_{artt'} =$

$$\begin{cases} t-t'+1 & \text{if } t' \le t < t' + CTime_{ar} - 1 \\ CTime_{ar} & \text{if } t \ge t' + CTime_{ar} - 1 \end{cases}$$

The right-hand-side indicates the cumulative number of plane-days available for type a aircraft up to day t. The left-hand-side accounts for all possible plane activities up to day t, whether flying or inventoried. The inventory terms are straightforward. The delivery and recovery terms work as follows: if a delivery initiated on day t' is completed by the end of day t, then the entire time $Ctime_{ar}$ (which may be integer or fractional) is included in the left-hand-side of the cumulative balance constraint for day t. On the other hand, if a delivery initiated on day t' is not completed by the end

of day t, then only the time expended so far, $t - t' + 1$, is counted in the day t constraint.

An experiment attesting to the value of the cumulative aircraft balance constraints is reported in [Morton, Rosenthal, and Lim, 1996]. If the $Ctime_{ar}$'s were all integer, these constraints would be redundant and could be omitted.

3.6.3 Aircraft Capacity Constraints

Troop Carriage Capacity Constraints:

$$TPAX_{uart} \leq MaxPAX_a \cdot X_{uart} \quad \forall u, a, r, t : t \in T_{uar}$$

Maximum Payload Constraints:

$$TonsUE_{uart} + PAXWt \cdot TPAX_{uart} \leq MaxLoad_{ar} \cdot X_{uart} \quad \forall u, a, r, t : t \in T_{uar}$$

Cargo Floor-Space Constraints:

$$PAXSqFt_a \cdot TPAX_{uart} + UESqFt_u \cdot TonsUE_{uart} \leq ACSqFt_a \cdot LoadEff_a \cdot X_{uart} \quad \forall u, a, r, t : t \in T_{uar}$$

3.6.4 Aircraft Utilization Constraints

The aircraft utilization constraints ensure that the total flying hours consumed by the fleets of each aircraft type over the planning horizon are within AMC's established utilization rates [Wilson, 1985; Gearing *et al.*, 1988]. These rates are meant to capture spares availability, aircraft reliability, crew availability, and other factors. The utilization constraints are formulated by comparing the flying hours consumed by an aircraft fleet in delivery and recovery flights to the maximum achievable flying hours for the fleet according to the utilization rate.

As an illustration of the above constraint, consider a fleet of five aircraft of the same type made available from day 11. If the utilization rate for this aircraft type is 10 flying hours per aircraft per day and the horizon is 30 days, then the maximum achievable flight time 1000 hours (10 hours/plane-day × 20 days × 5 planes). This total may not be exceeded for the whole

$$\sum_u \sum_{r \in R_a} \sum_{t \in T_{uar}} FltTime_{ar} \cdot X_{uart} + \sum_{r \in R_a} \sum_t FltTime_{ar} \cdot Y_{art} \leq \sum_t URate_a \cdot NPlanes_{at} \quad \forall a$$

fleet over the entire planning horizon, however, it is not unusual for a subset of aircraft to exceed utilization rates over a subset of the horizon, particularly during the early (surge) stage of a deployment.

3.6.5 Aircraft Handling Capacity at Airfields (MOG Constraint)

The aircraft handling constraints at airfields, commonly called MOG constraints, are perhaps the most difficult to model. This is because of two complicating factors that necessitate approximations. First, there is no airfield capacity data available that provides separate accounting of parking spaces and all the various services (refueling, maintenance, etc.). The MOG data provided by the Air Force is an approximation, attempting to aggregate all these services. Thus, the units of $MOGCap_{bt}$ are an idealized notion of airfield parking spaces (normalized to narrow-body sized aircraft), not a precisely defined physical quantity.

The second complicating factor in modeling airfield capacity is the congestion caused by the uncertainty of arrival times and ground times. A deterministic, time-discretized optimization model cannot accurately treat events occurring within a time period. For example, suppose the time period of the model is one day and an airfield has 20 landings per day. How much congestion occurs depends on when the landings occur during the day, a phenomenon not captured in the daily model. The MOG efficiency factor *MOGEff* is introduced to cushion the effect of not explicitly modeling uncertainty. In Section 6.2, we describe a stochastic programming model that more directly handles aircraft reliability and its effect on airfield capacity. The MOG constraints are formulated for each airfield and time period (as before, we use the notation $[Dtime_{abr}]$ to denote $Dtime_{abr}$ rounded to the nearest integer).

$$\sum_{u}\sum_{a}\sum_{r\in R_a}\sum_{\substack{t'\in T_{uar}\\ t'+[DTime_{abr}]=t}} (MOGReq_{ab}\cdot GTime_{abr}/24)\cdot X_{uart'} + \sum_{a}\sum_{r\in R_a}\sum_{t'+[DTime_{abr}]=t} (MOGReq_{ab}\cdot GTime_{abr}/24)\cdot Y_{art'} \leq MOGEff\cdot MOGCap_{bt} \quad \forall b,t$$

Dimensional analysis is useful for understanding these constraints. The right-hand-side is in the units of narrow-body parking spaces, because $MOGCap_{bt}$ is in those units and *MOGEff* is dimensionless. The first term on the left-hand-side accounts for airfield capacity consumed by all delivery missions that pass through airfield b during period t. The second term on the left does the same thing for recovery missions. The dimension of $MOGReq_{ab}$ is narrow-body parking spaces per plane, the dimension of $Gtime_{abr}/24$ is days, and the dimensions of $X_{uart'}$ and $Y_{art'}$ are planes per day; thus, the MOG constraints are dimensionally balanced.

Aircraft inventoried at origin or destination airfields do not consume any MOG capacity in the above formulation. This is not a mathematical limitation, but rather a modeling choice taken because inventoried planes do not consume ground services. It can be easily modified if data for "parking space MOG" and various "ground service MOGs" become available.

3.6.6 Initial Conditions

$$H_{ait}\equiv 0 \quad \forall a,i,t: t\leq 0$$
$$HP_{akt}\equiv 0 \quad \forall a,k,t: t\leq 0$$
$$X_{uart}\equiv 0 \quad \forall u,a,r,t: t\leq 0$$
$$Y_{art}\equiv 0 \quad \forall a,r,t: t\leq 0$$

4 FLEET-MIX TRADEOFF ANALYSIS

Prior to the C-17 Defense Acquisition Board (DAB) decision in November, 1995, there were a number of fleet options being considered as replacements for the aging C-141 fleet. These included "pure" C-17 fleets, as well as mixed fleets that included not only C-17s, but also a number of Non-Developmental Airlift Aircraft (NDAA), a Boeing 747-400F assigned the USAF designation C-33. THRUPUT II's first "operational" test supported the analysis required by the C-17 DAB.

Although many criteria must be considered when designing a fleet mix, a principal consideration is the ability to deliver the U.S. mobility requirements in support of our National Defense Strategy—currently two nearly-simultaneous Major Regional Contingencies (2-MRCs). Since THRUPUT II was designed to study strategic airlift, a two theater mobility study was a natural application of the model.

In the 2-MRC scenario, much of the cargo being flown from CONUS to the theaters is considered "out-sized" equipment, such as tanks or helicopters. Out-sized cargo is problematic, since it can only fit on certain wide-body aircraft, such as the C-5 or C-17. The C-33 is a hybrid in this regard; it can carry some, but not all types of out-sized cargo. THRUPUT II's features are well suited to contrast the capabilities of the long range, high payload C-33, with the more versatile, but smaller C-17. It was conceivable that THRUPUT II would show the lifting capability of a modest C-33 fleet could move most of the bulk and over-size cargo, allowing C-5s to satisfy the out-size requirement. Alternatively, the results might show that the demand for out-sized cargo movement dominates, and that the C-5 must be supplemented with C-17s to meet that requirement.

An additional fleet mix tradeoff involves the consumption of ground resources. The C-17 is designed to onload and offload quickly in an austere environment, while the C-33 is principally an airliner, and requires longer runways and a more robust support infrastructure. However, unless refueled in flight, the C-17 needs to stop more frequently than a C-33, which could offset any advantage derived from its reduced ground requirements. These two contrasting aspects of C-17 and C-33 resource utilization could interplay so as to give one aircraft considerable advantage in a contingency.

Cargo loading and airfield utilization are just two of a myriad of issues surrounding the procurement of any new airlifter. Without detailed modeling and simulation, the C-17 DAB

could not hope to make an informed choice based on objective criteria. However, unlike previous boards, this time the analysis included results provided by a detailed optimization model.

4.1 Input

THRUPUT II's input requirements are generalized into four categories: 1) unit, 2) airfield, 3) aircraft, and 4) route data. The source of the unit movement requirements is called the Time-Phased Force Deployment Data, or TPFDD. This highly detailed list of equipment and personnel requirements is intended to identify everything necessary to carry out our national strategy. Consequently, it can be quite detailed and extremely long. In fact, the TPFDD used in this analysis initially consisted of more than 21,000 entries. Modeling each of these entries as a THRUPUT II unit was unthinkable, given current computational limitations. Through careful screening and consolidation (see Section 5), the TPFDD was reduced to just over 200 entries, each of which was read into THRUPUT II as a unit. From the pared TPFDD, we examined the origins (Aerial Ports of Embarkation—APOEs), and destinations (Aerial Ports of Debarkation—APODs) and attempted to set up a realistic enroute basing scheme that could support the movement.

The primary guidance for this airfield and TPFDD information was the Joint Chiefs of Staff, J8, force structure analysis called the Mobility Requirements Study, Bottom-Up Review Update, MRS-BURU [Joint Chiefs of Staff, 1995]. Its results were driven by a specific 2-MRC TPFDD, considered to be the most widely accepted requirements listing in existence. This report not only identified the who, what, when, and where of every movement requirement, but also listed the available airfields, including their relative capacities for air cargo traffic flow. MRS-BURU is credited for providing the motivation for upgrading the U.S. airlift fleet.

Compared with unit and airfield information, aircraft and route data were relatively straightforward to gather. Although an aircraft's effect on the airlift system is contentious, its performance characteristics are largely objective and easily derived. Route data presented a more difficult, yet not insurmountable challenge. Relying only on currently established AMC routing condemns the model to favor aircraft whose payload-range characteristics resemble the current fleet. Allowing THRUPUT II the latitude to choose new routes based on an aircraft's unique capabilities was preferable, so we offered many more route-aircraft combinations than might seem necessary at first glance. The tradeoff between making sufficient routes available and model tractability is discussed in [Toy, 1996].

In addition to the airlift system parameters, there were several subjective factors to consider when setting up the scenario. One such factor was the *MaxLate* parameter, which establishes how late cargo and passengers can arrive before incurring an extremely large nondelivery (no-go) penalty. Increasing *MaxLate* naturally allows more overall cargo to be delivered, but has the unfortunate effect of dramatically increasing the size of the model, since there are more feasible movement options. However, the need to keep the model small must be balanced with a reasonable estimate of when "late" becomes "too late" from an operational standpoint. For the purposes of this work, *MaxLate* was set at eight days, meaning any cargo or passengers that could not be moved by the Required Delivery Date (*RDD*) + 8 would be considered not delivered and cause the maximum penalty to be charged. Fortunately, the eight-day maximum affected all excursions similarly, thus mitigating any relative advantage attributable to this subjectivity when comparing fleet mixes.

As with all analyses, preparing the above inputs took vast amounts of time. However, we believe many of these inputs are not scenario specific, and can be re-used with little adjustment in a variety of studies.

4.2 Analysis

THRUPUT II's C-17/NDAA analysis was conducted parametrically by running each proposed fleet mix as a separate excursion. The performance of each fleet was evaluated primarily by how much of the movement requirement (cargo and passengers) was delivered in a timely manner. Examining unit "closure" in this way, we were able to identify several significant differences between fleets. One such difference is illustrated by Figures 1A and 1B. In these figures, days from the "kickoff" of the first contingency are given on the horizontal

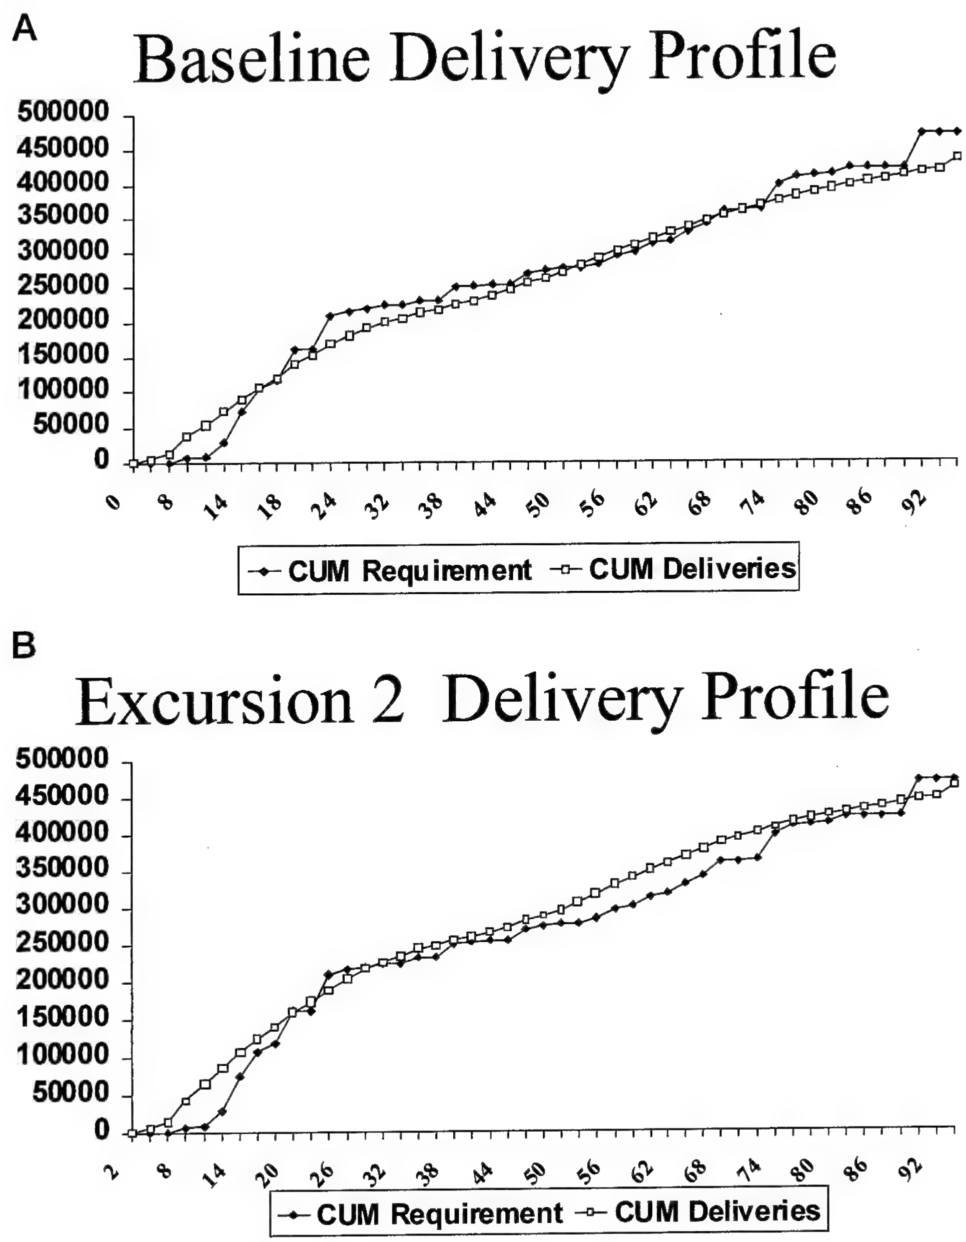

Figure 1. Tons delivered (vertical axis) versus contingency day (horizontal axis). The difference between the requirement and delivery lines in each scenario indicates that the aircraft fleet used in Excursion 2 performs better than the baseline.

axis. The two plots show the cumulative amount of cargo moved to date (stons), contrasted with the cumulative amount of cargo required to date. Figure 1A corresponds to the baseline fleet mix for the study. Figure 1B corresponds to one of the alternative fleet options under consideration.

Ideally, the airlift fleet should be able to accommodate the entire demand on time. However, given the non-uniform nature of the TPFDD requirements, all of the fleets examined had difficulty recovering from extremely large spikes in demand. The fleet used in the baseline case falls behind early and experiences great

difficulty catching up with the requirement. However, the fleet used in Excursion 2 experiences only brief lags in the cargo delivery. It was more able to move cargo early, and hence stayed ahead of the imminent demand surges.

Another key metric used to evaluate the different fleets was the relative proportions of on-time, late, and undelivered cargo and passengers. Figures 2A and 2B detail the results for the five cases examined. While all of the cases delivered similar amounts of on-time cargo, total cargo delivered (including late deliveries) varied significantly—notably between Excursion 1, and Excursions 2, 3, and 4. Interestingly, a comparison of the two figures shows that as the airlift fleet was tailored to improve cargo delivery, the number of passengers delivered went down. As a result, it appears that none of the proposed fleets dominate with respect to both cargo and passenger delivery.

During the course of our output analysis, we were unexpectedly enlightened by what began as a casual look at the marginals, or "shadow prices" associated with an optimal solution. Although not a key aspect of the study, airfield size played an enormous role in the overall performance of each of the fleets. The output revealed that relatively small changes in the airfield's capacity at key enroute and destination airfields would yield disproportionate changes in system performance. Moreover, these key bases differed depending on which fleet mix was under consideration. For example, fleets with many C-17s stressed enroute airfields considerably more than fleets with many C-33s. Conversely, excursions with a large number of C-33s relied heavily on destination airfield size, but did not require as extensive an enroute infrastructure due to their longer range. Given the performance characteristics of the two aircraft, this insight is not surprising, but the quantification of an airfield's marginal value is a great benefit of optimization that is unavailable in simulation. Moreover, this discovery clearly emphasized that a fleet mix decision is not one to be made in isolation. All aspects of the airlift system, including such factors as airfield infrastructure must be considered when choosing a mix of aircraft. They are not independent; treating them as such risks providing decision makers with skewed information about a critical piece of our nation's mobility force.

The analysis described here, performed by the THRUPUT II team in support of the 1995 C-17/NDAA DAB decision, is indicative of the type of insight that can be provided by this optimization model to a decision maker. We have elected to emphasize this theme rather than delve into the scenario and excursion specific details such as fleet composition and basing structure. However, one aspect of this project that does demand closer description involves the methods used to reduce such a large (indeed initially intractable) model to a manageable size.

5 MODEL REFINEMENTS FOR IMPLEMENTATION

5.1 Model Reduction

One of the key issues regarding implementation of optimization modeling, particularly in military applications, is the balancing of realism vs. tractability.

No mathematical model can ever be totally realistic. The optimization modeling process is itself a constrained optimization problem. The objective of the process is to maximize the amount of realism achieved, subject to the limitations on computational tractability. Regardless of the rapid rate of advances in computing, we will always be faced with finite limits on tractability and hence never achieve total realism. The question is: how much realism can one achieve with the resources at hand?

THRUPUT II has decision variables with as many as four indices, such as X_{uart}, so the crux of the balancing problem is the number of (u, a, r, t)-tuples included in any real instance of the model. In general, the more tuples allowed, the more realistic the model, but the more difficult it is to solve. The number of tuples depends on, first, how many of each index type exists (how many units, aircraft types, routes and time periods are modeled), and, second, what rules are used for allowing or prohibiting any (u, a, r, t) combination from being considered. These two aspects of model reduction are discussed next.

5.1.1 Aggregation

The number of units, aircraft types, routes and time periods in the instances we ran of the model were chosen with a great deal of attention to the issues raised above. Distasteful as it may seem, a certain amount of aggregation of

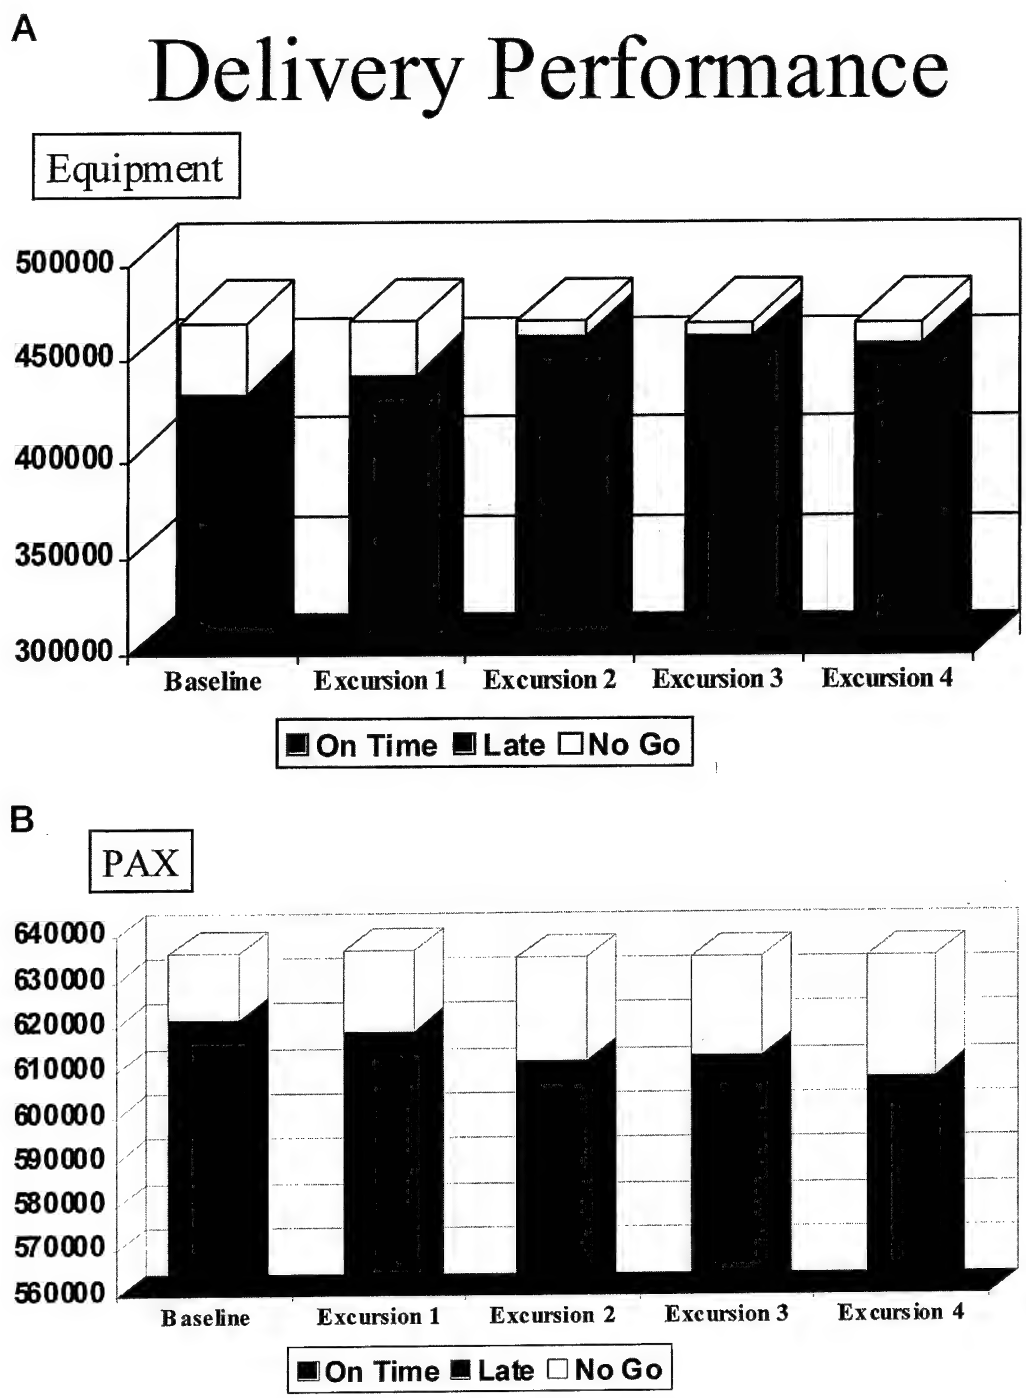

Figure 2. On time, late, and undelivered cargo for various airlift fleets.

entities is needed in any real-world modeling project. In this case, the most significant aggregation took place in the generation of airfields, units and time periods. Aggregation of airfields implies, in turn, a limit on the number of available routes.

We developed a location-theoretic optimization model for airfield aggregation, but in the case of the C-17/NDAA study, the USAF aviators on our team just used military judgement to decide which airfields to aggregate [Turker, 1995]. In the end, the infrastructure of the 2-MRC scenario was represented by 29 airfields.

The routes generated over the selected airfields were the product of a combination of the location-theoretic optimization model and the oversight of the aviators. The program used a tree structure to consider possible routes and screened them for inclusion based on various rules. The rules included: critical leg length of the aircraft, required crew rest or crew change, deviation of route length from great circle distance, aircraft/airfield compatibility (the civilian reserve fleet has landing restrictions not imposed on military aircraft and vice versa), and others. The oversight step was particularly intriguing, because the location and routing model was developed by a Turkish naval officer, who had no access to the real names and locations of the airfields during development and testing. Some routes had to be added or deleted based on understanding of the actual situation. The final result of this process was the inclusion of 313 routes for the entire scenario.

As stated in Section 4, the TPFDD file we were given for the C-17/NDAA analysis had over 21,000 movement requirements. This data set was first screened for the deletion of extremely small requirements. Then it was aggregated by assuming two movement requirements could be merged if the following conditions held: they had the same type of cargo (or passengers) to be moved, they had the same origin and destination (after airfield aggregations), and they had nearly simultaneous RDDs. The definition of "nearly simultaneous RDDs" was governed by a set of user-supplied parameters, which enforced simultaneity less rigorously as we went further out towards the horizon.

Aggregation of time is always a delicate issue in optimization modeling. Time has to be discretized, and nothing has a more direct effect on model size than the choice of time discretization units. For the C-17/NDAA study, we chose to divide time into two-day time periods, of which there were 47. Using 94 one-day time periods was intractable. See [Morton, Rosenthal, Lim, 1996] for an experiment on a single-MRC scenario, which showed that the cumulative aircraft balance constraints in Section 3.6.2 help lessen the effect of time discretization.

5.1.2 Variable Elimination and Sparsity

An algebraic modeling language such as GAMS is very conducive to implementing rules for limiting the number of admissible combinations of indices. In the case of the (u, a, r, t) tuples mentioned above, an X_{uart} variable is allowed to exist if all the following conditions hold:

- Route r flies from unit u's origin to its destination
- Aircraft type a can fly on route r with an acceptable payload
- The start time, t, of the mission is after unit u's available-to-load date
- The arrival time, if the mission starts at time t, is on or before $RDD(u) + MaxLate$
- There is a match between some cargo type (or passengers) that aircraft type a can carry and unit u's movement requirement
- Aircraft of type a must be available at the origin of route r at time t

These rules are evaluated once for each tuple and stored in a GAMS dynamic set. This set is referenced when the constraints are generated in order to achieve as much model sparsity as possible. The set is also used to eliminate other variables. For example, the aircraft inventory variables at destination airfields, HP_{akt}, cannot exist unless some aircraft of type a can potentially arrive at airfield k prior to time t.

There were other dynamic sets in THRUPUT II. In our experience with real-world optimization models, a serious investment of development time in the fine-tuning of dynamic sets for implementing model reduction rules can have a big payoff in tractability.

5.2 Computational Experience

After using the aggregations and reductions noted above, the model runs required for the C-17/NDAA study had the following problem dimensions: 200 units, seven (or fewer) aircraft types, 313 routes, 47 time periods and 29 airfields. The resulting model sizes and solution times are given below:

Scenario	Rows (000)	Cols (000)	Non-zeros	Solve Time
Baseline	161	183	1.9 mil	2.98 hrs
Ex 1	124	142	1.5 mil	1.95 hrs
Ex 2	154	177	1.9 mil	2.57 hrs
Ex 3	154	177	1.9 mil	3.14 hrs
Ex 4	154	177	1.9 mil	2.48 hrs

These runs were performed with GAMS as the problem generator and CPLEX 3.0 [CPLEX Optimization Inc., 1994] as the solver on an IBM RS6000/590 workstation. These rather large-scale optimization models presented a challenge. There were, in fact, several unsuccessful early attempts. We were very fortunate to be able to get advice from CPLEX Optimization, Inc. [Lowe, 1995] on solver settings. We sent them, via FTP over the Internet, a file containing a 3 million non-zero instance of the model, which they were able to solve. (This was when most but not all of the variable eliminations and sparsity refinements were implemented, so they have solved an even larger problem than the ones reported above.) The key advice from the CPLEX people was to use the barrier (interior point) algorithm, with tolerances and options tuned for this particular model.

5.3 Output

Each run of THRUPUT II for the C-17/NDAA study produced large amounts of output data. This profusion of information was too much for an analyst to absorb, so it had to be organized in relevant summary reports of the optimal solution. Since each case took a long time to solve, we made sure that all the optimal solution information was stored in readily accessed files. Then, a separate GAMS reporting program could be run many times against the same optimal solution. This proved to be useful because the AFSAA analysts often thought of new ideas for interesting summary reports while analyzing the old ones. Among the most widely used reports were:

- Total number of delivery missions flown, by aircraft type.
- Total number of missions flown, by route.
- For each unit, the closure date (date of last delivery if unit is fully delivered) displayed next to the ALD and RDD, along with the total amounts delivered on-time, late and not at all.
- For each MRC and time period, cumulative deliveries vs. requirements, separated by bulk, over-size, out-size and passengers.
- For each airfield, a report on MOG utilization, summarized as the number of days when MOG use exceeds P percent of capacity for P = 10, 25, 50, 75, 90, 95, and 100.

6 MODEL EXTENSIONS

6.1 Solution by Cascade

Although the difficulties associated with solving a large THRUPUT II model can be partially redressed by the model reduction techniques just described, ongoing research at the Naval Postgraduate School demonstrates that THRUPUT II may be solved in a piecemeal manner, thus greatly increasing the allowable problem size. This section describes that effort.

Consider how a scheduler would approach the 2-MRC scenario. Meticulously optimizing all aircraft, loading, and route decisions over the entire scenario length is impossible for at least two reasons: 1) future uncertainty makes gathering accurate data for the latter periods of a scenario problematic, and 2) a sufficiently long contingency overloads the scheduler's ability to reconcile the myriad of decisions. A modeler formulating a linear program faces the same difficulties, namely incorporating the increasing problem size with decreasing certainty as the length of the scenario grows. For either scheduler or modeler, perhaps the most straightforward way of dealing with the difficulties incurred by a large scenario is to focus sequentially on a subset of the scenario's periods, then move forward in time to a new subset. This temporal "myopia" degrades the solution quality, but makes the problem simpler to solve. Moreover THRUPUT II, which is used to mimic scheduling but does not produce schedules, is more "accurate" if it can incorporate the

realism of nearsighted scheduling. For example, when choosing fleet size or infrastructure for use in future mobility contingencies, THRUPUT II ideally wishes to optimize given the *current scheduling capabilities*, instead of a utopian capability. A truly optimal schedule generated by THRUPUT II might alter decisions made at the outset of a contingency based on specific delivery requirements several weeks later. This is unrealistic, and can be avoided by reducing the ability of the formulation to look so far ahead.

The *proximal cascade* heuristic applied to THRUPUT II proceeds by solving for all variables and constraints whose domain is defined for the first 20 (for example) periods. Thus missions are flown so as to minimize delivery penalties in the first 20 periods, subject to the constraints applicable in those periods. Then, the process is cascaded forward in time to solve for a later set of periods. Mathematically, this implies generating a feasible solution by successively solving for only a subset of rows and columns, then moving to a set of rows and columns corresponding to later time periods. Each of these subproblems should overlap the previously solved subproblem in order to minimize the end effects caused by the former's temporal limitation. Fortunately, this methodology is facilitated by the structure of THRUPUT II. Variables and constraints in this model directly affect only nearby time periods. For example, missions flown on day 5 of a scenario have a large impact on the missions that can be flown on day 7, but only a minor impact on the missions that can be flown on day 25. This characteristic manifests itself as an overlapping "staircase" along the main diagonal of an LP's constraint coefficient matrix. The width of the overlap gives the number of time periods directly affected by the decisions (variable levels) made in a given time period. The rest of the coefficient matrix is relatively sparse, since variables (columns) associated with the early time periods rarely appear in constraints (rows) corresponding to the later time periods. This well known methodology is known as either the rolling horizon, or *proximal cascade* heuristic. However, the heuristic is sparsely documented, and is theoretically incomplete, since no scheme to bound the solution quality has been offered.

The quality of the solution produced by the *proximal cascade* heuristic is dependent on many scenario specific factors, and cannot be stated theoretically for most problems. However, a bound on the solution quality may be derived by exploiting information derived from this heuristic solution. Since a given time period is only directly linked to a few adjacent time periods, relaxing the rows associated with these nearby periods can separate subproblems out of an otherwise linked model. As with most decompositions however, the success of this scheme is dependent on the ability to compute accurate prices for resource consumption of the relaxed constraints. With such prices, a Lagrangian penalty can be applied to the subproblems, and a lower bound can be derived. Often, price selection is computationally intensive, which makes Lagrangian methods undesirable. However, in this case, reasonable prices are readily available from the *proximal cascade* heuristic just computed.

The *proximal cascade* heuristic offers a way to produce a more realistic schedule than an unencumbered optimization model. It also greatly reduces the tractability problems associated with the large models demanded by mobility planners. Finally, the cost of scheduling myopia may be estimated by solving a series of relaxed subproblems. For these reasons, the method shows great promise for use with THRUPUT II [Baker, 1997].

6.2 Incorporating Aircraft Reliability

Aircraft reliability is an important factor in the ability of the airlift system to deliver troops and materiel in a timely fashion. The current fleet has a mix of planes with differing reliability characteristics. For example, the C-5 requires unscheduled maintenance on approximately 15% of its landings while the rate for the newer (and smaller) C-17 fleet is under 7% (1994 peace-time data, AMC). Broken aircraft reduce the lift capability of the system by reducing the size of the effective fleet. In addition, aircraft requiring unscheduled maintenance and repairs reduce throughput by consuming scarce resources (e.g., maintenance, crew-duty hours, ramp space) that might otherwise contribute to on-time deliveries.

Simulation models for airlift systems are attractive because they can incorporate high levels of detail such as tracking individual aircraft and incorporating unscheduled maintenance and repairs. However, simulation models typically use naive aircraft routing and

scheduling rules; as a result, it is possible to provide a simulation model with additional resources (e.g., more aircraft or routing options) and yet have system performance degrade. Linear programming models use more aggregate representations of the airlift fleet and infrastructure and do not incorporate uncertainty. However, due to optimal scheduling and routing, linear programming models better lend themselves to analysis of system bottlenecks by providing marginal values on specific resources, and in some cases, LP models may be more appropriate for comparing system performance under different sets of resources.

A *stochastic optimization* model for strategic airlift combines the ability of a simulation to include uncertain aircraft ground times with an LP's ability to optimally schedule and route aircraft. However, the resulting stochastic optimization model is typically very large and requires special-purpose optimization software. We have extended the LP model of Section 3 to incorporate aircraft reliability [Goggins, 1995]. The model is identical to the deterministic model except that the ground time $GTime_{abr}$ which appears in the airfield capacity constraint (Section 3.6.5) is replaced with a discrete random variable and the modified constraint includes an elastic decision variable which allows the constraint to be violated at a certain cost.

Mathematically, these modifications can be summarized as follows. Let ω denote a specific ground-time scenario (such as a scenario where a C-5A breaks, and its repair time is seven hours), and let p_{bt}^{ω} be the probability of observing scenario ω for a particular base b and time t. $GTime_{abr}^{\omega}$ represents the "effective ground-time" spent by aircraft a at base b when flying route r under ground time scenario ω, and $MOGPen_{bt}$ is the unit penalty for violating the airfield capacity at base b in time t. The elastic decision variable $R_{bt}^{\omega} \geq 0$ denotes the amount by which capacity is exceeded at airfield b in time t under scenario ω. The new airfield capacity constraints and the additional objective function term are specified in the following two equations.

The stochastic optimization model has been solved for the modest-sized data set in [Lim, 1994] which has 20 units, seven aircraft types, 17 airfields, and 30 time periods. Three of the seven aircraft were modeled as having random ground times (C-5, C-17, and C-141) and we assumed each aircraft type breaks independently. Ground times were approximated by discrete distributions with 9 realizations for each aircraft type, resulting in $9^3 = 729$ realizations for each base b and time t combination. The resulting stochastic model increases the number of airfield capacity constraints by a factor of 729 over the deterministic model from $30 \times 17 = 510$ to 371790. There are an equal number of additional decision variables of type R_{bt}^{ω}. We solved the stochastic model with a Benders' decomposition algorithm. While the total number of constraints in the stochastic model is greater by a factor of more than 50, the increase in running time over the deterministic model is a factor of 12 (20 minutes to 100 seconds on an IBM RS6000 590 workstation) [Goggins, 1995].

Airfield Capacity Constraint when Aircraft have Random Ground Times:

$$\sum_u \sum_a \sum_{r \in R_a} \sum_{t' + [DTime_{abr}] = t} (MOGReq_{ab} \cdot GTime_{abr}^{\omega}/24) \cdot X_{uart'}$$
$$+ \sum_a \sum_{r \in R_a} \sum_{t' + [DTime_{abr}] = t} (MOGReq_{ab} \cdot GTime_{abr}^{\omega}/24) \cdot Y_{art'} - R_{bt}^{\omega}$$
$$\leq MOGCap_{bt} \quad \forall b, t, \omega$$

Additional Objective Function Term to Discourage Capacity Violations:

$$\sum_b \sum_t \sum_{\omega} p_{bt}^{\omega} \cdot MOGPen_{bt} \cdot R_{bt}^{\omega}$$

Since the linear and stochastic programming models contain more aggregate representations of the airlift system, we examined whether the schedules proposed by these optimization models were "flyable" in a more detailed simulation. We developed a discrete-event stochastic simulation model that took as input the output of these mathematical programs; specifically, the simulation model attempts to execute a proposed aircraft routing schedule. The strategy of coupling optimization and simulation models in this way is very attractive. Confidence can be gained in the optimization model as certain parameters are tuned (as we describe below) and the performance of the simulation can be improved since naive

scheduling rules are replaced with those proposed by an optimization model.

Our experimental results on the modest-sized data set of [Lim, 1994] compared schedules proposed by the linear and stochastic programming models. During the peak demand periods, we observed a 10% increase in cargo and troop deliveries when the simulation model executed schedules proposed by the stochastic program. Because the stochastic optimization model is larger and more difficult to solve (and cannot currently be solved within algebraic modeling languages such as GAMS), it is desirable to "tune" the deterministic optimization model so that it yields delivery schedules that are achievable in the simulation model. As described in Section 3.6.5, controlling the "MOG efficiency value" *MOGEff* is one way to achieve this. We empirically determined that a *MOGEff* value of 0.80 gave deterministic optimization schedules that were "flyable" in the stochastic simulation.

7 CONCLUSIONS AND ONGOING RESEARCH

THRUPUT II is an optimization model of the airlift mobility system that has proven useful to Air Force analysts in an important acquisition study. The Air Force analytical community has in the past put much more reliance on simulation than on optimization. This is in contrast to civilian industries, such as petroleum, electronics, airlines, forestry and many others, where optimization is very widely used.

While we were developing THRUPUT II, a similar and concurrent effort was under way at the RAND Corporation. The CONOP model of [Killingsworth and Melody, 1994] is also a GAMS-based, multi-period linear programming model for airlift optimization. It has some features not found in THRUPUT II. In May, 1996, the NPS and RAND groups started a joint effort to develop a new optimization model with the best features of both THRUPUT II and CONOP. The new model is called the NPS/RAND Mobility Optimizer (NRMO). Among NRMO's features that are not modeled in THRUPUT II are:

- The use of tankers for aerial refueling, and the facility for some tankers to change roles between refueling and cargo hauling.
- The modeling of shuttle flights and ground transportation in theater: some units have the option of direct delivery vs. transshipment, and some aircraft have the option of changing roles between strategic carriers and shuttlers.
- Detailed flow balance and utilization constraints for crews.
- The modeling of recovery bases, so that aircraft arriving in theater have the option of receiving services and crew changes at some other airfield besides the MRC's main port of debarkation.

The NRMO model is currently in use in a study of airfield infrastructure and in a large Pacific scenario. The detailed formulation of NRMO and the results of these studies will be given in a future report [Melody *et al.*, 1996].

REFERENCES

Baker, S., *A Cascade Approach for Staircase Linear Programs with an Application to Air Force Mobility Optimization*, PhD Dissertation, Operations Research Department, Naval Postgraduate School, Monterey, CA., June 1997.

Brooke, A., Kendrick, D., and Meeraus, A., *GAMS A User's Guide*, The Scientific Press, 1992.

CPLEX Optimization Inc., *CPLEX*, version 3.0 for *GAMS*, Incline Village NV, 1994.

Gearing, R., Hill, J., and Wilson, D., "UTE Rates Revisited," *Airlift*, Spring 1988.

Goggins, D., *Stochastic Modeling for Airlift Mobility*, M.S. Thesis, Operations Research Department, Naval Postgraduate School, Monterey, CA., September 1995.

Joint Chiefs of Staff, "Mobility Requirements Study—Bottom-Up Review Update" (U) SECRET, The Pentagon, Washington, DC., March 1995.

Killingsworth, P., and Melody, L., "CONOP Air Mobility Optimization Model," RAND Corp., Santa Monica CA, July 1994.

Lim Teo Weng, *Strategic Airlift Assets Optimization Model*, M.S. Thesis, Operations Research Department, Naval Postgraduate School, Monterey, CA., September 1994.

Lowe, T., Telephone interviews and electronic mail, CPLEX Optimization Inc., Incline Village NV, September 16–18, 1995.

Melody, L., Rosenthal, R., Morton, D., and Baker, S., "NPS/RAND Mobility Optimizer (NRMO)," Fifth Air Force Mobility M&S Users' Group, 1–2 August 1996, Naval Postgraduate School, Monterey, CA.

Merrill, D., Electronic mail, HQ Air Mobility Command Studies and Analyses Group, Scott AFB IL, September 29, 1997.

Morton, D., Rosenthal, R., and Lim Teo Weng, "Optimization Modeling for Airlift Mobility," *Military Operations Research*, 1, No. 4, Winter 1996.

Schubert, K., Whisman, A., and Steppe, J., "Airfield Capability Modeling," 64th MORS Symposium, 18–20 June 1996, Fort Leavenworth, KS.

Stucker, J., "Airfield Capacity Estimator (ACE)," Fifth Air Force Mobility M\&S Users' Group, 1–2 August 1996, Naval Postgraduate School, Monterey, CA.

Toy, A., *Route Prioritization for an Airlift Mobility Optimization Model*, M.S. Thesis, Operations Research Department, Naval Postgraduate School, Monterey, CA, September 1996.

Turker, Y., *Route and Column Generation Methods for Airlift Mobility Optimization*, M.S. Thesis, Operations Research Department, Naval Postgraduate School, Monterey, CA, December 1995.

Wilson, D., "UTE: Utilization Rate: What Is It, How Is It Derived and How Should It Be Used?," *Airlift*, Winter 1985.

Wing, V., Rice, R., Sherwood, R., and Rosenthal, R., "Determining the Optimal Mobility Mix," Joint Chiefs of Staff (J8), The Pentagon, Washington, DC, October 1991.

Yost, K., "The THRUPUT Strategic Airlift Flow Optimization Model," Air Force Studies and Analyses Agency, June 1994.

ERRATUM: BENCHMARKING AND EFFICIENT PROGRAM DELIVERY FOR THE DEPARTMENT OF DEFENSE'S BUSINESS-LIKE ACTIVITIES

In the article "Benchmarking and Efficient Program Delivery for the Department of Defense's Business-Like Activities" that appeared in the Spring 1996 issue of *Military Operations Research* there is an error in the DEA model (2) presented on page 24. The formulation as it appears in the article is as follows:

$$\text{maximize } h_0 = \frac{u_s S_0}{v_{wy}WY_0 + v_{sqft}SQFT_0 + v_{apf}APF_0}$$

$$\text{subject to: } \frac{u_s S_1}{v_{wy}WY_1 + v_{sqft}SQFT_1 + v_{apf}APF_1} + \frac{u_s S_2}{v_{wy}WY_2 + v_{sqft}SQFT_2 + v_{apf}APF_2}$$

$$+ \ldots \frac{u_s S_{237}}{v_{wy}WY_{237} + v_{sqft}SQFT_{237} + v_{apf}APF_{237}} \leq 1$$

$$\frac{u_s}{v_{wy}WY_0 + v_{sqft}SQFT_0 + v_{apf}APF_0} \geq \epsilon$$

$$\frac{v_{wy}}{v_{wy}WY_0 + v_{sqft}SQFT_0 + v_{apf}APF_0} \geq \epsilon$$

$$\frac{v_{sqft}}{v_{wy}WY_0 + v_{sqft}SQFT_0 + v_{apf}APF_0} \geq \epsilon$$

$$\frac{v_{apf}}{v_{wy}WY_0 + v_{sqft}SQFT_0 + v_{apf}APF_0} \geq \epsilon$$

This formulation shows each ratio of weighted outputs to weighted inputs to be summed and that the sum must be less than or equal to one. This is incorrect. The correct formulation is to have each ratio to be less than or equal to one. Following is the corrected model (2) as it should have appeared in the article.

$$\text{maximize } h_0 = \frac{u_s S_0}{v_{wy}WY_0 + v_{sqft}SQFT_0 + v_{apf}APF_0}$$

$$\text{subject to: } \frac{u_s S_1}{v_{wy}WY_1 + v_{sqft}SQFT_1 + v_{apf}APF_1} \leq 1$$

$$\frac{u_s S_2}{v_{wy}WY_2 + v_{sqft}SQFT_2 + v_{apf}APF_2} \leq 1 \vdots \frac{u_s S_{237}}{v_{wy}WY_{237} + v_{sqft}SQFT_{237} + v_{apf}APF_{237}} \leq 1$$

$$\frac{u_s}{v_{wy}WY_0 + v_{sqft}SQFT_0 + v_{apf}APF_0} \geq \epsilon$$

$$\frac{v_{wy}}{v_{wy}WY_0 + v_{sqft}SQFT_0 + v_{apf}APF_0} \geq \epsilon$$

$$\frac{v_{sqft}}{v_{wy}WY_0 + v_{sqft}SQFT_0 + v_{apf}APF_0} \geq \epsilon$$

$$\frac{v_{apf}}{v_{wy}WY_0 + v_{sqft}SQFT_0 + v_{apf}APF_0} \geq \epsilon$$

In the corrected model there would be 241 constraints—one for each commissary (237 commissaries) plus one for each output and input variable. In the incorrect model that appeared in the article there would have been only five constraints.

I would like to note that the error did not affect the results of the research presented in the article. The model as used in the actual analysis was correct. Only the presentation of the model in the paper was incorrect.

I would like to thank Major Greg Hoscheit of the Army Recruiting Command for noticing the error and pointing it out.

About our Authors

AN INTEGER SOLUTION HEURISTIC FOR THE ARSENAL EXCHANGE MODEL (AEM)

by Daniel J. Green, James T. Moore, and John J. Borsi

Dan Green is an Air Force captain working on his PhD in Operations Research at the Air Force Institute of Technology (AFIT). Before coming to AFIT, Dan served on missile combat crew at Whiteman AFB, MO. His next assignment is to the Air Force Logistics Management Agency at Maxwell AFB, AL. Jim Moore is an Air Force lieutenant colonel and an Associate Professor of Operations Research at the Air Force Institute of Technology. His research interests and teaching include integer programming, large scale optimization, and applications of operations research to military problems. John Borsi is an Air Force lieutenant colonel at Air Mobility Command (AMC) where he is a military analyst. Before arriving at AMC, John was an Assistant Professor of Operations Research at the Air Force Institute of Technology. His research interests include linear programming and air mobility modeling and analysis.

ON THE PRINCIPLE OF "FALSE RANGING" IN WEAPONRY

by Israel David

Israel David teaches probability, statistics and OR in the university of the Negev (desert region of the holyland). His current research interests concern applications of stochastic dynamic programming to manufacturing systems and to the services. Dr. David spent a decade (1976–1986) as a Military OR analyst for the IDF, first in the Artillery Hq. and then in the joint Ground Forces Hq.

RISKNAV™: A DECISION AID FOR PRIORITIZING, DISPLAYING, AND TRACKING PROGRAM RISK

by C. C. Cho, P. R. Garvey, and R. J. Giallombardo

Mr. Garvey is Senior Principal Scientist in the Economic and Decision Analysis Center at The MITRE Corporation, where he leads MITRE's risk management and analysis work program. His work program is supporting several Air Force sponsors with computer-based analytical decision-aids for program risk management, as well as for theater missile defense. Mr. Garvey is widely published in the area of risk management and analysis. He is a four-time recipient of the Best Paper award from the Department of Defense Cost Analysis Symposium for his work in this area. His technical articles have appeared in publications by IEEE Computer Society Press, Springer-Verlag's Economics and Mathematical Systems series, and in professional journals such as *Military Operations Research* and *The Journal of Cost Analysis*. He is presently completing a textbook *"Probability Methods for Cost Uncertainty Analysis—A Systems Engineering Perspective"*, to be published in 1998 by Marcel Dekker, Inc. Mr. Garvey completed his undergraduate and graduate degrees in mathematics and applied mathematics from Boston College and Northeastern University.

Mr. Giallombardo is a Group Leader in the Economic and Decision Analysis Center at the MITRE Corporation. He is responsible for work program direction and resource allocation in support of several DOD sponsors. His current research interests include information technology economics, software resource analysis and modeling, risk analysis, health economics, and econometrics. Mr. Giallombardo holds an MS in Business Economics from Bentley College Graduate School of Business and a BA in Physics from Colby College. He is a member of Phi Beta Kappa and is a Society of Cost Estimating and Analysis (SCEA) Certified Cost Estimator/Analyst.

Mr. Chien-Ching Cho is Principal Operations Research Analyst in the Econmic and Decision Analysis Center at The MITRE Corporation. He co-chairs MITRE's Operations Research Cluster Group and has supported numerous projects for the Air Force's Electronic Systems Center. His technical strengths include operations research techniques, applied statistics, and cost/schedule risk analysis. His current research interests are process modeling, software cost/schedule estimation, and weapon-target assignment. Mr. Cho holds a B.S. from Tunghai University, Taiwan, an M.S. from State University of N.Y. at Buffalo, and Ph.D. from University of Wisconsin-Madison, all in Industrial Engineering. He is a member of the Military Operations Research Society.

THE PROBABILISTIC MULTIPLE-TRAVELLING-SALESMEN FACILITY-LOCATION PROBLEM: SPACE-FILLING CURVES AND ASYMPTOTIC EUCLIDEAN ANALYSES

by Yupo Chan and David L. Merrill

Yupo Chan has 25 years of postdoctoral experience in industry, universities, and government, including an honorary assignment as a Congressional Fellow. His research interests center around transportation science, networks and combinatorial optimization, multicriteria decision-making facility-location and spatial temporal information. Dr. Chan earned all his three degrees from MIT. He swims seven days a week, which explains why his hair is bleached grayer and grayer every time you see him.

David Merrill is senior analyst at the Analysis Group of the US Air Force Air Mobility Command at Scott Air Force Base. He is a retired career Air Force officer with extensive flying experience. His recent policy-oriented analyses include the acquisition of the Air Force C17 fleet. Among his professional honors is the Koopman Prize of the Institute of Operations Research and Management, awarded to the subject paper in 1991.

APPLICATION AND EXTENSION OF THE THRUPUT II OPTIMIZATION MODEL FOR AIRLIFT MOBILITY

by Richard E. Rosenthal, Steven F. Baker, Lim Teo Weng, David F. Fuller, David Goggins, Ayhan 0. Toy, Yasin Turker, David Horton, Daniel Briand, and David P. Morton

This paper involved the efforts of many individuals at the Naval Postgraduate School. Professors Richard Rosenthal and David Norton (now at UT Austin) served as the team leaders. The student members were: Steven Baker (USAF), Lim Teo Weng (Republic of Singapore Air Force), David Fuller and David Goggins (USN), and Ayhan Toy and Yasin Turker (Turkish Navy).

The remaining team members were David Horton and Daniel Briand (USAF), who supervised the research team from AFSAA at the Pentagon.

Although everyone participating in this research has been reassigned, many of the members continue to work on airlift mobility issues.

Editorial Policy and Submission of Papers

EDITORIAL POLICY

The title of our journal is *Military Operations Research.* We are interested in publishing articles that describe *operations research* (OR) methodologies used in important *military* applications. We specifically invite papers that are significant military OR applications. Of particular interest are papers that present case studies showing innovative OR applications, apply OR to major policy issues, introduce interesting new problem areas, highlight educational issues, and document the history of military OR. Papers should be readable with a level of mathematics appropriate for a master's program in OR.

All submissions must include a statement of the major contribution. For applications articles, authors are requested to submit a **letter** to the editor—exerpts to be published with the paper—from a **senior decision-maker** (government or industry) stating the benefits received from the analysis described in the paper.

To facilitate the review process, authors are requested to categorize their articles by **application area** and **OR method**, as described in Table 1. Additional categories may be added. (We use the MORS working groups as our applications areas and our list of methodologies are those typically taught in most OR graduate programs.)

INSTRUCTIONS TO *MILITARY OPERATIONS RESEARCH* AUTHORS

The purpose of the "instructions to *Military Operations Research* authors" is to expedite the review and publication process. If you have any questions, please contact Ms. Corrina Ross, MORS Communications Manager (email: morscm@aol.com).

General

Authors should submit their manuscripts (3 copies) to:

Dr. Gregory S. Parnell, Editor, *Military Operations Research*
Military Operations Research Society
101 South Whiting Street, Suite 202
Alexandria, VA 22304

Per the editorial policy, please provide:

- authors statement of contribution (briefly describe the major contribution of the article)
- letter from senior decision-maker (application articles only)
- military OR application area(s)
- OR methodology (ies)

TABLE 1: APPLICATION AREAS & OR METHODS

Composite Group	APPLICATION AREA	OR METHODOLOGY
I. STRATEGIC & DEFENSE	Strategic Operations Nuclear Biological Chemical Defense Arms Control & Proliferation Air & Missile Defense	**Deterministic Operations Research** Dynamic Programming Inventory Linear Programming Multiobjective Optimization Network Methods Nonlinear Programming
II. SPACE/C41SR	Operational Contribution of Space Systems C41SR Operations Research & Intelligence Information Warfare Electronic Warfare & Countermeasures Unmanned Systems Military Environmental Factors	**Probabilistic Operations Research** Decision Analysis Markov Processes Reliability Simulation Stochastic Processes Queuing Theory
III. JOINT WARFARE	Land & Expeditionary Warfare Littoral Warfare/Regional Sea Control Power Projection, Planning, & Execution Air Combat Analysis & Combat ID Special Ops/Operations other than War Joint Campaign Analysis	
IV. RESOURCES	Mobility & Transport of Forces Logistics, Reliability, & Maintainability Manpower & Personnel	**Applied Statistics** Categorical Data Analysis Forecasting/Time Series Multivariate Analysis Neural Networks Nonparametric Statistics Pattern Recognition Response Surface Methodology
V. READINESS & TRAINING	Readiness Analytical Support to Training & Mission Rehearsal Battlefield Performance, Casualty Sustainment, & Medical Planning	
VI. ACQUISITION	Measures of Effectiveness Test & Evaluation Analysis of Alternatives Cost Analysis Decision Analysis	
VI. ADVANCES IN MILITARY OR	Modeling, Simulation, & Gaming Revolution in Military Affairs (Long Range/Strategic Planning) Computing Advances in Military OR	**Others** Advanced Computing Advanced Distributed Systems (DIS) Cost Analysis Wargaming

Length of Papers

Submissions will normally range from 10–30 pages (double spaced, 12 pitch, including illustrations). Exceptions will be made for applications articles submitted with a senior decision-maker letter signed by the Secretary of Defense.

Paper Electronic Submission with Figures, Graphs and Charts

After the article is accepted for publication, an electronic version of the manuscript must be submitted in Microsoft Word or WordPerfect. For each figure, graph, and chart, please include a camera-ready copy on a separate page. The figures, graphs, and tables should be of sufficient size for the reproduced letters and numbers to be legible. Each illustration must have a caption and a number.

Mathematical and Symbolic Expressions

Authors should put mathematical and symbolic expressions in Microsoft Word or WordPerfect equations. Lengthy expressions should be avoided.

Approval of Release

All submissions must be unclassified and be accompanied by release statements where appropriate. By submitting a paper for review, an author certifies that the manuscript has been cleared for publication, is not copyrighted, has not been accepted for publication in any other publication, and is not under review elsewhere. All authors will be required to sign a copyright agreement with MORS.

Abbreviations and Acronyms

Abbreviations and acronyms (A&A) must be identified at their first appearance in the text. The abbreviation or acronym should follow in parentheses the first appearance of the full name. To help the general reader, authors should minimize their use of acronyms. If required, a list of acronyms can be included as an appendix.

Footnotes

We do not use footnotes. Parenthetical material may be incorporated into a notes section at the end of the text, before the acknowledgment and references sections. Notes are designated by a superscript letter at the end of the sentence.

Acknowledgments

If used, this section will appear before the references.

References

References should appear at the end of the paper. The references should be unnumbered and listed in alphabetical order by the name of the first author.

Appendixes

If used, this section will appear after the reference.

MILITARY OPERATIONS RESEARCH SOCIETY, INC.

COPYRIGHT ASSIGNMENT AGREEMENT

THIS AGREEMENT, made by and between ______________________, hereinafter ("Assignor(s)") and MILITARY OPERATIONS RESEARCH SOCIETY, INC., hereinafter Assignee).

Assignor(s) (*for U.S. Government employees: to the extent legally permissible*) hereby grants, transfers, assigns, and conveys to Assignee, its successors and assigns, the entire title, right, interest ownership and all subsidiary rights in and to an article entitled ____ (hereinafter the "Work"), including, but not limited to, the right to secure copyright registration of the Work and to any resulting registration in Assignee's name as claimant. Notwithstanding the foregoing, Assignor(s) reserves all proprietary rights other than copyright, such as patent rights; a royalty-free right to reproduce the Work; and the right to prepare derivative works based upon the copyright in the Work.

Assignor(s) hereby confirms that Assignor(s) owns the entire title, right and interest in the work, including the right to reproduce, prepare derivative works based upon the copyright in the Work, and distribute the Work, whether or not the Work constitutes a "work made for hire" as defined in 17 U.S.C. Section 201(b); Assignor(s) agrees that no rights in the Work are retained by Assignor(s) except as stated above. Assignor(s) agrees to execute any documents that might be necessary to perfect Assignee's ownership of copyrights in the Work and to registration.

Assignor(s) represents that the Work has not been copyrighted or published; that it is not being submitted for publication elsewhere; and, if the Work is officially sponsored by the U.S. Government, that it has been approved for open publication.

This Agreement constitutes the entire agreement between the parties hereto; this Agreement supersedes any prior oral or written agreement or understanding between the parties; and, **in the case of a Work Made for Hire, this Agreement has been signed by the Assignor(s)'s employer.** This Agreement shall only be effective if and upon the date that the Work is accepted by Assignee for publication in *Military Operations Research*, the Journal of the Military Operations Research Society or PHALANX, the Bulletin of the Military Operations Research Society.

ASSIGNOR(S)

By: ______________________

Title: ______________________

Date: ______________________

ASSIGNOR(S)

By: ______________________

Title: ______________________

Date: ______________________

MILITARY OPERATIONS RESEARCH SOCIETY, INC. (ASSIGNEE)

By: ______________________

Title: ______________________

Date: ______________________

MORS MILITARY OPERATIONS RESEARCH SOCIETY

Subscriptions, Monographs & Workshop Reports Order Form

SHIP TO:	Return form to:
Name: ______	Military Operations Research Society
Organization: ______	101 S. Whiting Street, #202
Address: ______	Alexandria, VA 22304-3418
______	Phone: 703-751-7290
Phone: ______	FAX: 703-751-8171
All orders must include a phone number.	E-mail: morsoffice@aol.com

Please note: **All prices are for US orders only! For INTERNATIONAL prices please call the MORS office.**

PHALANX Bulletin subscription (Published Quarterly)	1 year—$20	2 years—$35		
Military Operations Research (MOR) Journal subscription	1 year—$40	2 years—$75		
SAVE$$** Combined both subscriptions and **SAVE$$ *PHALANX* and *MOR* Journal	**1 year—$55**	**2 years—$100**		

DESCRIPTION		QTY	Freight	PRICE	TOTAL
Simulation Technology 1997, Sessions 1,2 & 3	(1988)			$ 10.00	
Human Behavior and Performance as Essential Ingredients in Realistic Modeling of Combat—MORIMOC II, Vol 1 & 2	(1989)			$ 55.00	
MORIMOC III	(1990)			$ 10.00	
MORS/ITEA *Emphasizing the "E" in T&E*	(1991)			$ 10.00	
Simulation Validation Workshop—SIMVAL II	(1992)			$ 10.00	
Emphasizing the Cost in COEA—MORS/SCEA Mini-symposium	(1993)			$ 10.00	
ITEA/MORS *How Much Testing is Enough?*	(1994)			$ 10.00	
Joint Requirements Oversight Council Process Workshop (JROC)	(1994)			$ 10.00	
Simulation Data and Its Management Workshop (SIMDATAM)	(1993 & 1995)			$ 10.00	
Quick Response Analysis Methodology (QRAM)	(1996)			$ 15.00	
MOR Analyst's Handbook, Vol I: *Terrain, Unit Movement, and Environment*			$ 5.00	$ 20.00	
MOR Analyst's Handbook, Vol II, Area 1: *Conventional Weapons Effects (Ground)*			$ 3.00	$ 5.00	
SAVE$$ Combine *Military OR Analyst's Handbook,* Vol I and Vol II, Area 1			$ 5.00	$ 25.00	
Warfare Modeling			$ 5.00	$ 35.00	
Military Modeling for Decision Making, Third Edition			$ 5.00	$40.00	
Sales TAX applicable on BOOKS delivered in Virginia only! Sales tax DOES NOT apply to freight or subscriptions.			**VA Residents**	**4.5% TAX**	
❑ CASH ❑ CHECK ❑ VISA ❑ MC ❑ AMEX			TOTAL ENCLOSED		$

Credit card # ______

Expiration date: ______ Billing zip code: ______

Name on Credit Card (print): ______

Signature: ______

TOTAL $

CALL THE MORS OFFICE FOR CREDIT CARD ORDERS

THE MILITARY OPERATIONS RESEARCH SOCIETY SYMPOSIUM (MORSS)

Currently, an average of over 950 military operations research analysts attend the annual Military Operations Research Society Symposium (MORSS), a national professional society under the sponsorship of the Joint Staff, US Army, US Navy, US Air Force, US Marine Corps and the Office of the Secretary of Defense. The attendees represent the services, the private sector, academe, and other government agencies. The MORS Symposium provides them with a unique opportunity to hear the thoughts of influential military leaders. It also provides them with the opportunity to exchange information and examine completed or ongoing research in a classified setting.

The symposium consists of a series of meetings centered around 32 working groups. The working groups also meet in composite sessions to address a wider spectrum of topics, which are of interest to their associated composite group. In addition, MORS offers general sessions and tutorials. The general/special sessions include thematically related invited papers of broad interest, selected best working group papers, workshop reports and sessions, and an education session. The tutorials are generally educational sessions on tools and techniques of operations research.

The 66th MORS Symposium will be held at the Naval Postgraduate School (NPS) in Monterey, California on 23-25 June 1998. This year's theme, *Preparing for Military Operations Research in the 21st Century*, was echoed in remarks recently made by the Chief of Naval Operations concerning graduate education at NPS: "Students will expand their breadth of knowledge in a particular discipline and will reinvigorate their ability to successfully analyze and solve the complex challenges we face. These important skills will help guide our Navy into the 21st Century through fresh thinking and innovation." Because of its mission, the Naval Postgraduate School is the perfect setting for our society to prepare a path to the next century.

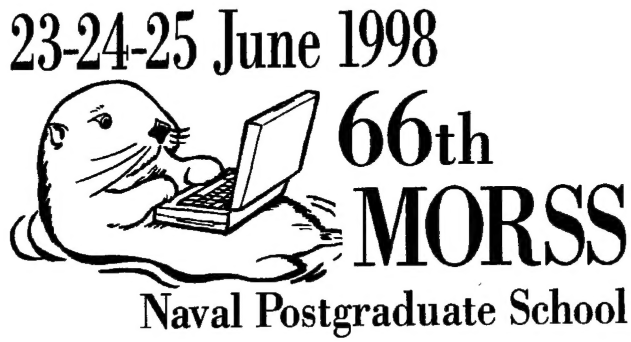

This year's symposium will consist of a keynote session on Tuesday, three general sessions, eight composite group/working group sessions, three tutorial sessions during lunch, two poster sessions, and the social event Wednesday evening.

So, mark your calendars now for the 66th Military Operations Research Society Symposium at the historic Naval Postgraduate School from 23-25 June 1998. Get ready to enjoy one of the most beautiful spots on the face of the earth — the Monterey Peninsula. Help us provide a path to the next century by contributing to a successful symposium with an early response to the Call for Papers and submitting your "approved for public release" abstract in January. And while you're convincing your boss to let you go to the Symposium, convince him or her to go with you! If you have any doubts at all on presenting a paper, have any other questions, have not received the ACP, or would like to assist the program staff, please contact the MORS office at (703) 751-7290 or CDR Kirk Michealson at (703) 697-0064.

MILITARY OPERATIONS RESEARCH SOCIETY
BOARD OF DIRECTORS

OFFICERS

President
Dr. Jerry A. Kotchka*
(314) 777-9990
E-mail: jkotchka@mdc.com

Vice President Finance and Management
Dennis R. Baer*
(703) 312-2149
E-mail: drbaer@aol.com

Vice President Meeting Operations
Susan M. Iwanski*
(516) 346-9138
E-mail: sue_iwanski@atdc.northgrum.com

Vice President Professional Affairs
Dr. Robert S. Sheldon*
(703) 648-8268
E-mail: bs@s3i.com

Secretary
CAPT Lawrence L. Dick*
(703) 602-2791
E-mail: leesmail@pipeline.com

Past President
Frederick E. Hartman*
(202) 298-7166
E-mail: hartmanfe@aol.com

Executive Vice President
Richard I. Wiles*
(703) 751-7290
E-mail: morsone@aol.com

Vice President Administration
Natalie S. Addison
(703) 751-7290
E-mail: morsvpa@aol.com

*** - Denotes Executive Council Member**

OTHER DIRECTORS

Col Thomas Allen, USAF, AFSAA
LTC(P) James E. Armstrong, USA, AD*, USMA
Mary T. Bonnet, AFSAA
Dr. Alfred G. Brandstein, MCCDC
Dr. Yupo Chan, AFIT
Dorn Crawford, AD, US ACDA
Dr. Henry C. Dubin, USA OPTEC
James B. Duff, PRC, Inc
Helaine G. Elderkin, FS, AD, Computer Sciences Corporation
Brian D. Engler, Systems Planning and Analysis, Inc.
Christine A. Fossett, AD, GAO
Maj Mark A. Gallagher, USAF, USSTRATCOM
Priscilla A. Glasow, AD, MITRE
Dr. Dean S. Hartley, Data Systems R&D Program
Dr. Glen H. Johnson, US ACDA
RADM Pierce J. Johnson, USNR, Naval Reserve Readiness Command
COL James L. Kays, USA, USMA
Kerry A. Kelly, AD, USSTRATCOM
COL Kenneth Konwin, USAF, DMSO
MAJ Willie J. McFadden, USA
CDR Kirk Michealson, USN, OSD PA&E
Dr. Julian I. Palmore, University of Illinois
Dr. Gregory S. Parnell, FS, AD, Virginia Commonwealth University
Anne M. Patenaude, SAIC
Dr. Roy Rice, Teledyne Brown Engineering
Dr. Patricia A. Sanders, OSD PA&E
Edward A. Smyth, The Johns Hopkins University
Dr. Cyruss Staniec, Logicon
Dr. Stuart H. Starr, AD, MITRE
Joseph A. Tatman, AD, TASC
Harry J. Thie, AD, RAND
LCDR Katie P. Thurman, USN, XO, NRD
Howard G. Whitley III, USA CAA
James L. Wilmeth III, AD, SETA Corp.
LTC Mark A. Youngren, USA, AD, Naval Postgraduate School

*** - Denotes Advisory Director**